arXiv:1210.7375v2 [physics.soc-ph] 30 Oct 2012

TRACTABLE AND CONSISTENT RANDOM GRAPH MODELS

ARUN G. CHANDRASEKHAR[‡] AND MATTHEW O. JACKSON[⋆]

Abstract. We define a general class of network formation models, Statistical Exponential Random Graph Models (SERGMs), that nest standard exponential random graph models (ERGMs) as a special case. We analyze conditions for practical and consistent estimation of the network formation parameters. This addresses two holes in the estimation of exponential random graph models. First, although it is known that due to the enormity of the space of possible networks Markov chain Monte Carlo methods of estimation face slow (exponential) mixing times for many specifications, we provide a first set of results identifying nontrivial specifications for which practical, accurate estimation is possible. Second, we provide consistency results showing when maximum likelihood and GMM (generalized method of moments) estimates of parameter converge to the true values in SERGMs. In particular, we show that a class of SERGMs that count subgraphs of various types are consistently estimated using direct and fast techniques. We also define a related class of network formation models, SUGMs, and show that they are also consistently and easily estimated when networks are sufficiently sparse. We illustrate the application of the models and techniques with data.

JEL Classification Codes: D85, C51, C01, Z13.

Keywords: Random Networks, Random Graphs, Exponential Random Graph Models, Exponential Family, Social Networks, Network Formation, Consistency, Sparse Networks

1. Introduction

> ...[A] pertinent form of statistical treatment would be one which deals with social configurations as wholes, and not with single series of facts, more or less artificially separated from the total picture.
>
> *Jacob Levy Moreno and Helen Hall Jennings, 1938.*

Social networks shape our opinions and behaviors and so it is essential to have models for the estimation of network formation. Since an individual's decision to form a relationship with another individual can depend on whether they have friends in common, how many other relationships they are each involved in, and other aspects of their social positions, relationships are generally far from independent. This is seen in

Date: Revision: October 2012.

We thank Larry Blume for a discussion of the paper at the January 2012 ASSA meetings. We also thank Ben Golub for detailed comments. Chandrasekhar thanks the NSF Graduate Research Fellowship Program. Jackson gratefully acknowledges financial support from the NSF under grants SES-0961481 and SES-1155302 and from grant FA9550-12-1-0411 from the U.S. Air Force Office of Scientific Research and the Defense Advanced Research Projects Agency.

[‡]Microsoft Research New England; Department of Economics, Stanford University.

[⋆]Department of Economics, Stanford University; Santa Fe Institute; and CIFAR.

many settings social and economic settings: from informal favor exchange where the presence of friends in common can facilitate robust favor exchange (e.g., Jackson et al. (2012)), to international trade agreements where the presence of one trade agreement can influence the formation of another (e.g., Furusawa and Konishi (2007)). Similarly, in forming a network of contacts in the context of a labor market, an individual benefits from relationships with others who are better-connected and hence relationships are not independently distributed (e.g., Calvo-Armengol (2004)); nor are they in a setting of risk-sharing (e.g., Bramoullé and Kranton (2007)).

Once such interdependencies exist, estimation is no longer at the level of pairs of nodes, but must encompass more of the network, even the whole network, as reflected in the quote from Moreno and Jennings (1938) above. Given their ability to incorporate general dependencies, exponential random graph models (henceforth "ERGMs") have become the workhorse models for estimating network formation.[1] Indeed, as originally shown via a powerful theorem by Hammersley and Clifford (1971), the exponential form can nest *any* random graph model and can incorporate arbitrary interdependencies in connections.[2] Moreover, there are random utility foundations for these models, as we show and others have shown.

Although such models are widely used and are seemingly natural tools for estimating the formation of networks with interdependent ties, there are two critical gaps in the understanding of these models of network formation and we address both of them in this paper.

First, the number of possible networks on a given number of nodes is an exponential function of the number of nodes. Estimating the likelihood of a given network (as part of Bayesian, maximum likelihood, GMM, or other procedures) requires having some estimate of its likelihood relative to the other networks that could have appeared instead. Given the exponential explosion of the number of possible networks,[3] it is impossible to directly calculate likelihoods of a given network appearing and so some approximation is necessary. However, this computational hurdle is so formidable that even state-of-the-art algorithms for estimation can be inadequate and inaccurate. This is the subject of a burgeoning literature that shows that some MCMC estimation techniques cannot mix in less than exponential time for large classes of exponential random graph models (e.g., see Bhamidi et al. (2008) and Chatterjee et al. (2010)).[4]

[1]These grew from work on what were known as Markov models (e.g., Frank and Strauss (1986)) or $p*$ models (e.g., Wasserman and Pattison (1996)). An alternative approach is to simply work with regression models at the link (dyadic) level, but to allow for dependent error terms, as in the "MRQAP" approach (e.g., see Krackhardt (1988)). Although that approach can be quite clean econometrically, it is less well suited for identifying the incidence of particular network subgraph structures that may be implied by various social or economic theories.

[2]Their theorem applies to undirected and unweighted networks. See the discussion in Jackson (2008). Of course, the representation can become fairly complicated; but the point is that the ERGM model class is broadly encompassing. Furthermore, as we illustrate below, it can also be adapted to allow for multigraphs such that nodes can have multiple types of relationships.

[3]For example with only $n = 30$ nodes, the number of possible networks is 2^{435} (which exceeds most estimates of the number of atoms in the universe).

[4]Apart from estimation issues, there is another issue that concerns which formulations of exponential random graph models are distinguished from independent link models. For example, Chatterjee and Diaconis (2011) show that some classes of ERGMs are indistinguishable in a well-defined sense from independent-link models. Our formulations avoid such issues, and in particular our results on sparse

Second, there is little that is known about the consistency of estimators in such exponential random graph models.[5] It is important to emphasize that the relevant asymptotic frame is one in which the number of nodes grows to infinity, rather than the number of networks; as in many applications, the data consist of a single network.[6] If links were all independent of each other, then a single network provides many observations and link probabilities are estimable. However, once there are interdependencies between links, standard asymptotic techniques do not apply and that is why so little is known about the consistency of estimation techniques in this setting. This does not mean that consistency is precluded, (just as it is not in time series or spatial settings) as even though the probability must be specified at the network level there is still a lot of information that can be discerned from the observation of a single network.[7] Nonetheless, it does mean that we must approach the asymptotics in a way that accounts for the potentially complex correlations and interdependencies that arise in link formation.[8]

To fix ideas and discuss our approach in more detail, let us detail the issues that ERGMs face. In an ERGM, the probability of a network of observing a particular network g with an associated vector of statistics $S(g)$ (for example, density of links, number of cliques of given sizes, the average distance between nodes with various characteristics, counts of nodes with various degrees, etc.) is proportional to

$$\exp\left(\theta \cdot S\left(g\right)\right)$$

where θ is a vector of model parameters.[9] Turning this into a probability requires normalizing this expression so that the probability of observing g is

$$\mathrm{P}_\theta\left(g\right) = \frac{\exp\left(\theta \cdot S\left(g\right)\right)}{\sum_{g'} \exp\left(\theta \cdot S\left(g'\right)\right)}. \tag{1.1}$$

networks apply to classes of models that are well-distinguished from independent-link models and to which the Chatterjee and Diaconis (2011) results do not apply.

[5]Of course, consistency has been examined in the context of some random graph models (e.g., see Bickel et al. (2011); but not for the class of models that we consider here. Consistency of a different sort has been examined by Shalizi and Rinaldo (2012) in the context of ERGMs: whether or not if one observes only a subsample of the network, whether the estimated parameters will accurately reflect the true parameters of the full network. That is also related to work by **?**, but we do not address that separate issue here.

[6]If the asymptotic frame was one in which we observed a growing number of networks with a fixed number of nodes, each generated by the same network formation process, consistency would be trivial. Although there exist such data sets, they are the exception rather than the rule. And, even for such data sets, the estimation issues that we tackle here are relevant.

[7]In cases where a time series of link formation is observed or postulated and estimated, then one can take advantage of the sequentiality to see how each link forms conditional on the network in place at the time (e.g., see Christakis et al. (2010)). However, without such information, or in cases where links may further evolve over time, the network perspective again prevails. For example, Mele (2011) considers a sequential model and then shows that it becomes effectively equivalent to a certain ERGM.

[8]Clearly, there are other settings with interdependencies, such as time series and spatial settings. The network setting presents a complex set of interdependencies that do not permit off-the-shelf approaches.

[9]The reasons for using the exponential family is clear. It nests many standard distributions such as multivariate normal, Poisson, power, lognormal, exponential, gamma, beta, Weibull, Laplace, multinomial, etc. Also, it has many nice properties regarding its cumulants and moment generators, and has been well-studied in the statistics literature.

The challenge is that estimating the parameters of such a model, by maximum likelihood, Bayesian, GMM, or other methods, requires estimating how the relative likelihood of a network relates to parameters. This involves either explicitly or implicitly estimating the denominator of (1.1). However, as we mentioned above, examining all possible networks g' is infeasible even for small n, and thus, one has to take other routes.

The adaptation of Markov chain Monte Carlo (MCMC) sampling techniques to sample networks and estimate ERGMs, by Snijders (2002) and Handcock (2003), provided a breakthrough and the subsequent development of computer programs based on those techniques led to their widespread use.[10] However, it was clear to the developers and practitioners that at least some specifications of ERGMs had convergence problems. Until recently, it remained unknown if or when these techniques would mix accurately in a feasible time. Given the enormity of the space any MCMC procedure would visit only an infinitesimal portion of the space, and it was unclear whether that was sufficient to gain an accurate picture. Unfortunately, important recent papers have shown that for broad classes of ERGMs some of these MCMC procedures will take exponential time to mix unless the links in the network are approximately independent (e.g., see the discussions in Bhamidi et al. (2008) and Chatterjee et al. (2010)), in which case the ERGM was not needed as the network could then be viewed as a collection of independent edges and edge formation could be estimated directly. Such difficulties were well-known in practice to users of software that performs such estimations, as rerunning even simple models can lead to very different parameter and standard error estimates,but now are known to be more than an anomaly.

Beyond the computational challenges, it is also not known whether various estimators of ERGMs are consistent. For example, when they are computable, will maximum likelihood estimates converge to the true parameters as the number of nodes becomes large? Given that data in many settings consist of a single network or a handful of networks, we are interested in asymptotics where the number of nodes in a graph grows. However, it may be the case that increasing the number of nodes does not increase the information in the system. In fact, for some sequences of network statistics and parameters it is obvious that the parameters of the associated ERGM are *not* consistent. For example, suppose that $S(g)$ includes a count of the number of components in the network and the parameters are such that the network consists of a single or a few components. The limited number of components would not permit consistent estimate of the generative model. Thus, there are models where consistent estimation is precluded. On the other extreme where links are all independent, we know that consistent estimation would hold, and so the interesting question is for which models is it that consistent estimation can be obtained.

We do four things in this paper.

First, we propose a generalization of the class of ERGMs that we call SERGMs: *Statistical ERGMs*. To understand the generalization, it helps to note that ERGMs define a probability distribution on the space of sufficient statistics $S(g)$ (e.g., frequencies of edges, cliques, stars, etc.) and conditional on having these characteristics, any

[10]See Snijders et al. (2006) for more discussion.

network is equally likely.[11] Thus, ERGMs the probability of forming a network is determined by its statistics; for instance, having a given link density, a given clustering coefficient, specific path lengths, etc., and then any network exhibiting the same properties measured by the model is equally likely.[12] Moreover, the weighting (or *reference distribution*) used over the space of statistics is then a special one: the probability of seeing specific statistics depends on the number of networks that exhibit that statistic. It is also difficult to compute,[13] and it is not clear that it should be preferred by the modeler. SERGMs nest the usual ERGM models by noting that: (i) we can define the model directly over the statistics thus greatly reducing the dimensionality of the space, and (ii) we can weight the distribution over the space of statistics in many ways other than simply by how many networks exhibit the statistics. Some of these reference distributions generate natural models that both allow for realistic features as well as desirable statistical properties such as consistency and asymptotic normality of the estimators of model parameters.

Second, we examine sufficient conditions as well as some necessary conditions for consistent estimation of SERGM parameters (nesting ERGMs as a special case) and identify a class of SERGMs for which it is both easy to check consistency and estimate parameters. Models in this class are based on "count" statistics (for instance, how many links between nodes with certain characteristics exist, how many triangles[14] including certain types of nodes exist, how many nodes have a given degree, etc.). By slightly altering the reference distribution, we show that a natural family of models exhibit desirable statistical properties.

Third, we identify another class of network formation models that are based on the formation of subgraphs. We can think of such a network as being constructed from building blocks of varying sizes: links, triangles, larger cliques, stars, etc., layered upon each other. We show that if such models are sufficiently *sparse* in a well-defined way, then they are consistently estimable and parameters are asymptotically normally distributed. Such sparse networks appear in many applications as they have realistic features (e.g., average degree that grows at a rate less than n, but still allowing for high triadic closure, homophily, rich degree distributions, and so forth). These models are also easily to simulate from and admit random utility foundations that may be used in economic applications (e.g., counterfactual policy analysis, testing theory).

Fourth, we provide an illustration of the techniques developed here by applying them to data on social networks from Indian villages. In particular, we also develop an extension of the models that apply to multigraphs (so individuals may have different sorts of edges between them). This allows us to test several theories of multiplexing: fixed costs of link formation, patterns that foster favor exchange, and correlated (un)-observables. We find evidence consistent with predictions of a theory of favor exchange: being a member of a triangle may substitute for having multiple types of relationships with others.

[11]This is related to the well-known property of sufficient statistics of the exponential family.

[12]As an analogy, a binomial distribution defines the probability of seeing x heads but does not care about the exact sequence under which the x heads arrive.

[13]Just as an example, calculating the number of graphs on n nodes with a given number t of triangles is a challenging combinatorial problem.

[14]Triangles refer to triads: cliques of size three; that is, triplets of nodes i, j, k that include all three possible links.

In terms of related literature, our second class of models (SUGMs) is not the first to build networks up from randomly generated subgraphs, as that is done by Bollobás et al. (2011). Nonetheless, the results, models, and foci differ significantly. We are interested in building broad classes of models that have practical and simple estimators, and showing that those estimators are consistent. This involves a special technique in ordering the statistics in the way that they are counted, so, for instance, a link generated in isolation is counted differently from one that appears embedded in a clique. Bollobás et al. (2011) is more concerned with properties of the resulting networks such as thresholds for giant components, degree distributions, and on clustering properties, and thus looks at models with particular rates of subgraph formation.

The remainder of the paper is organized as follows. Section 2 introduces the framework, ERGMs, and discusses random utility foundations that are useful in economic applications. In Section 3 we introduce SERGMs. Consistent estimation of SERGMs is discussed in Section 4, and we also relate this to a weighted binomial distribution. In Section 5, we introduce another model of network formation based on subgraphs, which we show to be easily and directly estimable in the case of sparse networks, with consistent and asymptotically normally distributed estimators. We present simulation exercises as well as an empirical application to Indian village networks in Section 6, to show how well these models can mimic real-world network data. Section 7 describes extensions and applications of these models, including how they may be used in analyzing strategic network formation as well as multigraphs. We then provide an application to Indian village network data to study multiplexing between favor-exchange and other types of social links. Section 8 concludes.

2. Preliminaries and ERGMs

We begin with background definitions and notation.

2.1. Networks or Graphs.

Let $\mathcal{G}^n$ be a set of possible graphs on a finite number n of nodes. The class can consist of undirected or directed graphs and unweighted or weighted graphs. We take $\mathcal{G}^n$ to be finite, but it will generally be large as a function of n. For instance, if $\mathcal{G}^n$ is the set of all undirected, unweighted graphs on n nodes, then the cardinality of $\mathcal{G}^n$ is $|\mathcal{G}^n| = 2^{\binom{n}{2}}$.

We often omit notation $\mathcal{G}^n$ and denote a generic network by g. Thus, if we write $\sum_g$ it is understood to mean $\sum_{g \in \mathcal{G}^n}$ for whatever class of networks is relevant.

Unless otherwise stated we take $\mathcal{G}^n$ to be the set of undirected, unweighted graphs; but as will be clear, the results extend readily to more general classes, including multigraphs as we illustrate below.

2.2. Network Statistics.

Let $S(g) = (S_1(g), \ldots, S_k(g))$ be a vector of statistics concerning a graph $g \in \mathcal{G}^n$, where $S_l : \mathcal{G}^n \to \mathbb{R}$.

Given a statistic S, let

$$N_S(s) := |\{g \in \mathcal{G}^n : S(g) = s\}|$$

denote the number of graphs that result in a given statistic.

Statistics can be many different things, including standard network characteristics: the number of links, the number of cliques of some given size, the number of links that

have neighbors in common, homophily measures, numbers of nodes with given degrees, average distances, and so forth.

It is important to note that node characteristics can also be included in statistics. For example, nodes might be classified into some finite number of groups based on some characteristics, and then one can track the number of links between various types of nodes, various clustering and cohesion measures by types of nodes, and so forth. This permits the fitting of choice-based models, where the utility that an individual derives from a link to another depends on node characteristics and network position. See Section 7.1 for additional discussion.

2.3. **Standard ERGMs.** A standard ERGM model takes the form

$$\mathrm{P}_\theta(g) = \frac{\exp(\theta \cdot S(g))}{\sum_{g'} \exp(\theta \cdot S(g'))} = \frac{\exp(\theta \cdot S(g))}{\psi(\theta)} = \exp(\theta \cdot S(g) - \log(\psi(\theta))), \tag{2.1}$$

where $\theta = (\theta_1, ..., \theta_k) \in \mathbb{R}^k$ is a vector of real-valued parameters, and

$$\psi(\theta) = \sum_{g'} \exp(\theta \cdot S(g'))$$

is the appropriate normalizing coefficient.

2.4. **Maximum Likelihood Estimation.** Although our analysis has clear extensions to Bayesian estimation, we proceed mainly in the context of maximum likelihood estimation and of generalized method of moments.

Given a model of the form (2.1) and an observed network g, the maximum likelihood estimator, denoted $\widehat{\theta}(g)$, is the θ that maximizes $\mathrm{P}_\theta(g)$. Equivalently, it is the θ that maximizes $\theta \cdot S(g) - \log(\psi(\theta))$ and therefore satisfies

$$S(g) = \nabla\psi(\theta)/\psi(\theta),$$

and, by standard properties of the exponential family (see (2.2)), this can be written as

$$S(g) = \mathrm{E}_{\widehat{\theta}}[S(g')]. \tag{2.2}$$

We discuss estimation in more detail below.

2.5. **Asymptotics.** As mentioned in the introduction, although there are exceptions, many applications of interest involve observation of one (large) network. Given the interdependencies in links, this presents a challenge in estimation as there are not repeated independent observations. Thus, the challenge is understanding when a large enough network contains enough information to accurately estimate parameters. Correspondingly, our asymptotic frame is one in which we observe a single network and want to recover parameters consistently as the number of nodes goes to infinity. Therefore we discuss *sequences* of ERGMs, a sequence of distributions over $\mathcal{G}^n$, so that practical and consistent estimation of θ^n is possible for large enough n.

2.6. **Random Utility Rationalizations for ERGMs.** Given the extensive literature on strategic network formation, and the many applications where this is at least partly at play, it is important to be able to estimate strategic network formation models.

There are several ways in which ERGMs capture aspects of network formation that go beyond direct link values. One is that agents might value friends of friends, and thus wish to connect with others who have greater numbers of connections. As shown by Mele (2011) this leads to an ERGM specification that includes "two stars," or links Another is that agents might value bridging separate groups. Another is that agents might value having a relationship within a clique differently from having a relationship with someone with whom they do not share any friends.[15]

In Section 7.1 we detail specifications of SERGMs designed for strategic network formation.

3. SERGMs and Estimation Techniques

The expression for the normalizing coefficient of an ERGM, $\psi(\theta)$, is impossible to calculate directly with more than a few nodes, as the number of graphs on n nodes grows exponentially in n^2. This presents obvious problems for Bayesian, maximum likelihood estimation, and methods of moments estimations.

Our approach is based on re-specifying the model (to a more general class) that are based directly on the (sufficient) statistics $S(g)$ and then also possibly renormalizing the weights that this procedure generates.

3.1. An Example and Illustration of Computational Advantages.

We begin with an example that illustrates some of the ideas of our approach.

In many settings, the number of isolated nodes is significantly higher than one might observe at random in an Erdős-Rényi random graph. In fact, Moreno and Jennings (1938) estimated that there were more than ten million socially isolated individuals in the U.S. in the 1930's, many more times than would be expected if relationships were formed at random. This would be true in many settings in which individuals can decide not to be active and not to form any relationships, and is quite different from a setting where all nodes randomly form links and some just happen not to form any.

Let us thus consider a simplest model that allows nodes to be isolates, separately from happening not to form links.

In particular, let $S_L(g)$ denote the number of links in a network g and $S_I(g)$ denote the number of isolated nodes (with no links) under g. We can define a random network model where some nodes do not participate and then other nodes form links uniformly at random as in an Erdős-Rényi random graph.

The associated ERGM is given by

$$\mathrm{P}_\theta(g) = \frac{\exp(\theta_L S_L(g) + \theta_I S_I(g))}{\sum_{g'} \exp(\theta_L S_L(g') + \theta_I S_I(g'))} = \frac{\exp(\theta_L S_L(g) + \theta_I S_I(g))}{\psi(\theta_L, \theta_I)}.$$

[15]Valuing friends of friends relates to access to information and favors, as in the connections model of Jackson and Wolinsky (1996). Values to bridging appear in theories such as those underlying structural holes of Burt (1992). Values to cliques arise under various theories such as those of closure Coleman (1988) or favor exchange Jackson et al. (2012).

If $\theta_I = 0$ then this reduces to a standard Erdős-Rényi random network on all nodes. The more interesting case is where $\theta_I \neq 0$. In that case, networks become more ($\theta_I > 0$) or less ($\theta_I < 0$) likely based on the number of isolated nodes. For instance, if $\theta_I > 0$ then links would tend to be concentrated among a smaller number of nodes than under an Erdős-Rényi random graph, with a larger number of nodes with no links.

Many socially generated networks have excessive number of isolated nodes, and this is a simple model that generates networks with such effects.

The difficulty with estimating such a model is that the number of such networks in the calculation of $\psi(\theta_L, \theta_I)$ is $2^{\binom{n}{2}}$. With just 100 nodes this is 2^{4950}, while estimates of the number of atoms in the universe is less than 2^{258} (Schutz, 2003).

Methods of estimating the parameters of an ERGM are thus forced to circumvent direct calculation of $\psi(\theta_L, \theta_I)$, and instead use approximation methods such as MCMC (Markov Chain Monte Carlo) techniques.[16] Even with this approach, the fraction of all networks that can be sampled is tiny, and more pressingly, the space is difficult to traverse in a representative fashion. For instance, if one samples say 10000 networks, then one samples on the order of 2^{16} networks out of the possible 2^{4950} on 100 nodes. Thus, unless one is very lucky or very knowledgeable in choosing which networks to sample and how many to sample of different types, the sample is unlikely to be even remotely representative of the possible configurations that might occur. Indeed the time before which an MCMC technique has a chance to sample enough networks to gain a representative sample can be *exponential* in n^2 and so is prohibitively large even with a small number of nodes.[17] ,[18]

The alternative approach that we explore here takes advantage of the following insight.

Let $N_S(s_L, s_I)$ denote the number of networks that have exactly s_L links and s_I isolated nodes. An approximation of $N_S(s_L, s_I)$ is

$$K_S(s_L, s_I) = \begin{cases} \binom{n}{s_I}\binom{\binom{n-s_I}{2}}{s_L} & \text{if } \frac{n-s_I}{2} \leq s_L \leq \binom{n-s_I}{2}, \text{ and} \\ \\ 0 & \text{otherwise.} \end{cases}$$

The first expression $\binom{n}{s_I}$ is the number of ways in which one can choose s_I isolates out of n nodes, and the second expression $\binom{\binom{n-s_I}{2}}{s_L}$ is then the number of ways in which one can have s_L links among the remaining nodes. This is not a precise count, since $\binom{\binom{n-s_I}{2}}{s_L}$ allows for some isolated nodes among the nodes that are supposed to be connected. However, if $s_L \gg (n - S_I)\log(n - S_I)$, then for large n the number of networks among

[16] See Snijders (2002), Handcock (2003), as well as discussions in Snijders et al. (2006) and Jackson (2011). The rough intuition is that such methods sample some networks to get an impression of the relative likelihoods in terms of relative sizes of $\exp(\theta_L S_L(g') + \theta_I S_I(g'))$ and thus develop a rough estimate of the relative likelihood of the observed data under various specifications of θ.

[17] This does not even include difficulties of sampling. For example, as discussed by Snijders et al. (2006), a technique of randomly changing links based on conditional probabilities of links existing for given parameters can get stuck at complete, empty, or other extreme networks.

[18] Bhamidi et al. (2008) show that, in estimating many classes of ERGMs, MCMC techniques using Glauber or Metropolis-Hasting's dynamics (or any local Markov dynamic) mix in less than exponential time only if any finite group of edges are asymptotically independent.

the $\binom{\binom{n-s_I}{2}}{s_L}$ that end up having isolated nodes goes to 0.[19] Thus, in situations where expected degree exceeds the threshold of $\log(n - S_I)$, $K_S(s_L, s_I)$ becomes an accurate approximation to N_S. In fact, we further simplify by approximating K by using

$$\widetilde{K}_S(s_L, s_I) = \begin{cases} \binom{n}{s_I}\binom{\binom{n-s_I}{2}}{s_L} & \text{if } s_L \leq \binom{n-s_I}{2}, \text{ and} \\ \\ 0 & \text{otherwise,} \end{cases}$$

where the only difference from K is not imposing the restriction that $s_L \geq \frac{n-s_I}{2}$.

This is quite useful, since in settings where expected degree is above the threshold of $\log(n)$, we approximate

$$\psi(\theta_L, \theta_I) \approx \sum_{(s_L, s_I)} \widetilde{K}_S(s_L, s_I) \exp(\theta_L s_L + \theta_I s_I).$$

Here the summation involves less than $n(n(n-1)/2)$ entries, whereas the direct calculation involved $2^{n(n-1)/2}$ calculations. So, for instance, if $n = 100$ then this is fewer than 495000 calculations, as opposed to 2^{4950}. This makes the ERGM practically estimable, with a direct calculation of the likelihood.

We can then further simplify this expression to reduce the calculations further. The sum on the right hand side of (3.1) is concentrated in a shrinking neighborhood around the expected (and, with probability going to one, realized) values of S_I, S_L, so that one needs only sum in a shrinking neighborhood of the realized values of S_I, S_L (as we discuss in more detail below).

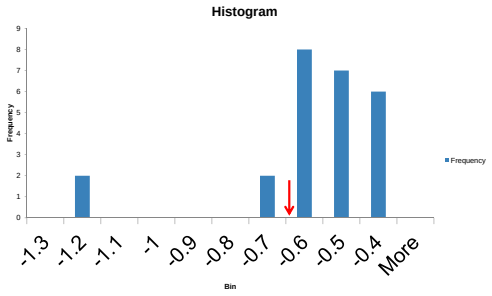
3.1.1. *An Example of the Problems with MCMC Estimation of ERGMs.* To illustrate this we simulate this model with $n = 30$. Of these nodes, 20 nodes form an Erdős-Rényi subgraph, with edges forming with i.i.d. probability of roughly $p = 1/3$ (the realized matrix had .34).[20] Meanwhile, the remaining 10 nodes are isolated. We fit a simple model using a standard ERGM estimation software (`statnet` via `R`) and estimate the following over 25 different runs on the same fixed network:

Figure 1 illustrates the potential instability associated with ERGM estimation using MCMC techniques. Even with only thirty nodes and a very simple two-statistic model, the algorithm for estimating the ERGM comes to very different estimates when rerun (and each MCMC run involves thousands of iterations). Even if an estimate provides for nontrivial standard errors on the estimates, a reliable algorithm should reproduce nearly the same estimates off of identical data when replicated. Given that this is a random sampling algorithm, some noise might be expected, but not of this magnitude.[21] Table 1 provides more details on the distribution of parameter estimates.

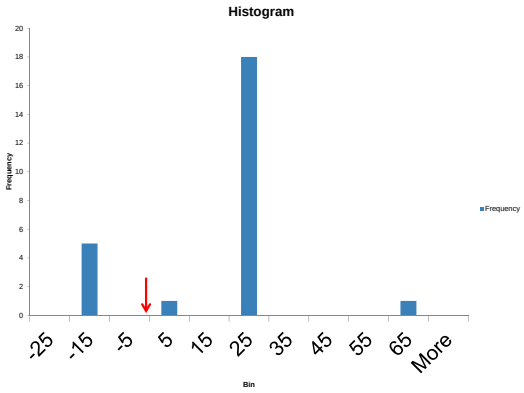
[19]This is a corollary to the fact that the threshold for connectivity in the "$G(n, M)$" class of random-graph models that Erdős-Rényi considered (the model where each possible network on n nodes having M links is given equal probability) is $\log(n)$ (see **?**). This implies that the fraction of networks with $M = \binom{n}{2}p$ links that have any isolated nodes goes to 0 above the threshold where $p = \log(n)/n$.

[20]This yields 190 observations of edge formation and 65 links were present and all nodes were connected.

[21]In fact, the parameters are well outside of the average standard errors that the program provides by a bootstrapping method, which are on the order of 0.11 for θ_L and 0.2 for θ_I. The standard errors in the program presume that the sample of draws in its estimation provide a good approximation for the population, and so the bootstrapping can be very inaccurate given that even thousands of iterations provides a completely inaccurate sample of the true population of networks and coefficients.



(A) statnet Estimates of the Links Parameter



(B) statnet Estimates of the Isolates Parameter

FIGURE 1. Histograms of estimated ERGM parameters by statnet across different runs of the same program on the same network. The left-hand figure is the histogram of estimated parameters on links, while the right-hand figure is the histogram of estimated parameters on isolates. The red arrow indicates the true parameters, which can also found directly using our estimation techniques as discussed below.

	$\widehat{\theta}_L$	$\widehat{\theta}_I$
Mean	-0.647	9.52
S.D.	0.1947	18.456
Median	-0.5992	15.8407
Min	-1.2564	-23.3957
Max	-0.4401	59.213

TABLE 1

The estimation with our SERGM technique leads to $\widehat{\theta}_L = -.65$ and $\widehat{\theta}_I = -.68$ (without any iterations as the likelihood is now a direct calculation).

It is worth noting that in the given model there is an even more direct and obvious estimation method. One can simply estimate p_I and p_L as the probabilities of being isolated and then links (conditional on not being chosen as an isolated node) by frequency counts. It follows that the associated θ are such that $\theta_j = \log \frac{p_j}{1-p_j}$ for $j \in \{I, L\}$. However, we have introduced the machinery of the $K(s)$-functions to preview a more general form of ERGMs that the researcher can study. In certain (sparse enough) cases, as discussed in Section 5, the even more direct estimators provide very fast and accurate estimates of this and many other specifications.

Before proceeding, we note that we used a model of isolates and links for illustrations, but there is nothing peculiar or special about this model in terms of the difficulties it poses for ERGMs nor for the improvements that SERGMs can offer in estimation.

3.2. **ERGMs in a Statistical Representation.**

As we saw from the previous example, it will generally be useful to express ERGMs in an alternative representation, taking advantage of the sufficient statistics of the realized networks. Instead of considering the probability of observing a particular *network*, we instead ask what the probability is of observing a particular realization of network *statistics*: for instance, what is the probability of observing a network with a given link density, clustering coefficient, number of cliques of some size, homophily index, and so forth. It is in the properties of the network that our theories lie.

More specifically, consider the probability of observing $S(g) = s$:

$$\mathrm{P}_\theta\left(S(g) = s\right) = \sum_{g:S(g)=s} \frac{\exp\left(\theta \cdot S(g)\right)}{\psi(\theta)}$$

and since $\psi(\theta) = \sum_{s'} N_S(s') \exp\left(\theta \cdot s'\right)$ it follows that

$$\mathrm{P}_\theta\left(S(g) = s\right) = \frac{N_S(s) \exp\left(\theta \cdot s\right)}{\sum_{s' \in A} N_S(s') \exp\left(\theta \cdot s'\right)}.$$

Generally, the range A of S is of much lower dimension than the set of all networks on n nodes. Thus, to the extent that one can compute/approximate $N_S(s)$, as we saw in the example given in Section 3.1 this can be a much easier formulation to work it. In fact, in many cases it makes sense to free up the N_S weights to a more general class, which can further simplify estimation and consistency.

3.3. **Statistical ERGMs (SERGMs).**

Note that (3.2) defines a model over the range A of potential network statistics of interest and, in principle, there is nothing special about the weighting $N_S(\cdot)$ function that is used over A. This leads us to our more general representation.

Consider a vector of network statistics $S = (S_1, \ldots, S_k)$ with potential values that lie in $A \subset \mathbb{R}^k$.[22] Any weighting function $K_S(\cdot)$ defined on A, together with a set of parameters $\beta \in \mathbb{R}^k$, define a *statistical exponential random graph model*, or SERGM for short. The associated probability of seeing realized statistics $S = s$ is:

$$\mathrm{P}_{\beta,K_S}\left(S = s\right) = \frac{K_S(s) \exp\left(\beta \cdot s\right)}{\sum_{s' \in A} K_S(s') \exp\left(\beta \cdot s'\right)} = \frac{K_S(s) \exp\left(\beta \cdot s\right)}{\phi(\beta)}. \tag{3.1}$$

where the normalizing coefficient is

$$\phi(\beta) = \sum_{s' \in A} K_S(s') \exp\left(\beta \cdot s'\right). \tag{3.2}$$

This is a model that states that the probability that a network exhibits a specific realization of statistics $S = s$ is given by an exponential function of the statistics s. Thus, the model is specified specifically on the properties of the network rather than the actual realized network. Given the realized statistics s, the network g such that $S(g) = s$ can be thought of as being drawn uniformly at random, but in fact that is inconsequential since the model is really about network properties rather than particular realized networks.

In the language of exponential families of random variables, $K_S(\cdot)$ is simply a *reference distribution*. Varying the reference distribution, of course, changes the resulting

[22]Given the finite number of possible networks, A is taken to be finite. If one wishes to work with weighted networks, then obvious extensions to continuous ranges and integrals apply.

odds of various values of s being drawn and can affect whether the model is consistently estimable. The special case in which $K_S = N_S$ is corresponds to an ERGM in a statistical representation.

Several remarks are in order.

First, note that there is no reason to maintain that $K_S(\cdot)$'s must approximate $N_S(\cdot)$'s. Nature may choose properties of networks (S's) according to some alternative weightings. The instance of studying $N_S(\cdot)$-weighted SERGMs may be a historical one: on another planet, people may have first modeled SERGMs and would see those as natural with the ERGMs being a special case.

Second, it can still be that one is interested in a specific sub-class of these models wherein the $K_S(\cdot)$'s are chosen to approximate, locally, the behavior of $N_S(\cdot)$'s, as we did in our isolates example. And, as we shall see, viewing things from the statistical perspective can help substantially with respect to estimation.

Third, one way to interpret this model is as follows. A distribution on properties $S \in A$ drives the network formation process, and those network properties are what the researcher ultimately cares about with respect to the social or economic theories being studied. Which network forms given the realized statistics is secondary and could be uniform at random, or according to some other conditional distribution, so long as given the realized s a network g such that $S(g) = s$ is drawn.[23]

3.4. **Estimation.**

Let us now examine maximum likelihood estimation for the general class of SERGMs, which nests the maximum likelihood estimation of traditional ERGMs that we saw in (2.2).

The maximum likelihood estimator $\widehat{\beta}(s)$ of a model of the form (3.1) is the maximizer of $\mathrm{P}_\beta(s)$, or equivalently:

$$\widehat{\beta}(s) = \operatorname*{argmax}_\beta \log\left(\mathrm{P}_\beta(s)\right) = \operatorname*{argmax}_\beta \beta \cdot s - \log(\phi(\beta)) + \log(K_S(s)).$$

Thus, except for extreme cases, it satisfies

$$s = \nabla\phi(\widehat{\beta}(s))/\phi(\widehat{\beta}(s)),$$

and so $s = \nabla\kappa_S(t)|_{t=0} = \mathrm{E}_{\widehat{\beta}(s)}[S]$, where κ is the cumulant. Therefore,

$$s = \frac{\sum_{s' \in A} K_S(s') \exp\left(\beta \cdot s'\right) s'}{\sum_{s' \in A} K_S(s') \exp\left(\beta \cdot s'\right)} = \mathrm{E}_\beta[S]. \tag{3.3}$$

[23]Note that any SERGM viewed as a distribution over statistics defined with weights K_S and parameters β also generates an ERGM, as it generates a distribution over networks, for example taking networks to be drawn uniformly at random from those with the given statistics. That ERGM can be written as

$$\mathrm{P}_{\beta,K}(g) = \frac{\frac{K_S(s)}{N_S(s)} \exp\left(\beta' s(g)\right)}{\sum_{g'} \frac{K_S(s(g'))}{N_{S(s(g'))}} \exp\left(\beta' s(g)\right)} = \frac{\exp\left(\beta' s(g) + f_n(s)\right)}{\sum_{g'} \exp\left(\beta' s(g) + f_n(s(g'))\right)}$$

where $f_n(s) = \log\left(\frac{K_S(s)}{N_S(s)}\right)$ Therefore the model is modified by a shift with weights $\log\left(K_S(s)/N_S(s)\right)$. The sum here is only over g' for which there is some realizable statistic value so that $N_S(s(g')) > 0$.

So, in a case where the SERGM is sufficiently identified, so that $\beta \neq \beta'$ implies that $\mathrm{E}_\beta[S] \neq \mathrm{E}_{\beta'}[S]$, then the maximum likelihood estimator $\widehat{\beta}(s)$ is a solution to

$$s = \mathrm{E}_{\widehat{\beta}(s)}[S], \tag{3.4}$$

provided that this is well-defined. For extreme values of s this will not be well-defined.[24] Here, we implicitly assume that the model is specified so that the probability of observing extreme statistics for which this is not satisfied is negligible, which will be true of the asymptotic specifications that we work with provided that the β's do not tend to extremes too quickly.

Of course, this also suggests that we may be interested in a GMM estimator associated with the first order condition equations normalized by the rate of growth of the associated network statistics S, or some other normalization. For example, we may have a diagonal matrix C with positive diagonal entries such that we are interested in the estimator

$$\widehat{\beta}(s) := \underset{\beta}{\operatorname{argmin}} \left(s - \mathrm{E}_\beta[S]\right)' C^2 \left(s - \mathrm{E}_\beta[S]\right). \tag{3.5}$$

The maximum likelihood estimator discussed above, $s = \mathrm{E}_\beta[S]$, then is the case where the min is achieved with a value of 0, whereas the above specification allows for the more general case.

4. Consistency

The next issue that we address is the consistency of estimation of SERGMs: Under what conditions does an estimator of a SERGM converge to the correct estimate in probability? To our knowledge there are no results on consistency of ERGM estimators. The primary challenge is that the data consists of a single network and the asymptotics are in terms of the number of nodes. Moreover, the interest in SERGMs is in the the fact that relationships are correlated, and so the data can be far from independent.

As with any estimation problem, for an estimator to be consistent, two things need to balance. First, the data have to have sufficient variation to be informative, i.e., the (norm of the) covariance of the data has to grow. That is, more information must be obtained from each additional unit of data (the change in statistics associated with each additional node). Second, the variation cannot grow too fast – that is, the rate of growth has to be such that, when appropriately normalized, the statistics will concentrate around their expected values so that the data induce the correct estimate of parameters. Both of these are usually guaranteed by having growing numbers of independent, or at least asymptotically independent (satisfying a mixing property) observations, so that new information comes in and a law of large numbers applies so that some normalized statistic concentrates around its mean. The challenge here

[24] For example, for a simple Erdos-Renyi random network where the count statistic is simply the number of links in the network, then if turns out that all links are present so that $s = n(n-1)/2$, then the β that maximizes the likelihood of the ERGM formulation is essentially infinite (the $\beta = \log\left(\frac{p}{1-p}\right)$ corresponding to the maximum likelihood estimator of the link probability ($p = 1$) is not well-defined). For more on the nonexistence of well-defined maximum likelihood estimates for extreme networks see Rinaldo et al. (2011) .

is that we cannot rely on approximate independence, and so must work directly with these conditions.

In our context, having sufficient variation rules out certain sorts of statistics. For instance, recall the example mentioned in the introduction such that the probability of a network depends on the number of components in the network. If the parameters are such that with high probability the network generated has only a few components, then there will be insufficient information in the data to accurately estimate the parameters of the model. We can thus see that statistics will typically have to take on a wide array of values in order to provide for accurate estimation, but that alone will be insufficient. The concentration about the mean will also be required. This requires that there not be too much interdependence between all the relationships in the network, so that asymptotically there is sufficient information for parameter estimation.[25]

We proceed as follows. First, we provide propositions that generally characterize consistency of SERGM estimation. Next, we then look at some specific SERGMs to show that they satisfy those conditions. Third, we also show that these classes of SERGMs can be easily (quickly) and accurately estimated.

4.1. **Consistency in Our Context.**

We begin with a general characterization of consistency of SERGMs, and then show how these conditions apply in specific contexts.

Consider a sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$ as defined in (3.1).

A sequence of SERGMs is *expectations-identified* with respect to a sequence of diagonal matrices $C_n > 0$ with positive diagonal entries if there exists $\gamma > 0$ such that

$$|C_n \mathrm{E}_\beta[S^n] - C_n \mathrm{E}_{\beta^n}[S^n]| > \gamma \, |\beta - \beta^n|$$

for all n.[26]

Expectations identification is an intuitive condition that requires that different parameters distinguish themselves with different means. It is a sort of minimal condition since if two different parameter values generate very similar expected statistics, then observing the realized statistic will not allow us to distinguish the parameters. It encodes, as shown in the appendix, a sufficient condition for the application of a uniform law of large numbers; thus, provided pointwise convergence is satisfied, one can show that parameter of an extremum estimator are consistent.

A sequence of SERGMs is *concentrated* with respect to a sequence of diagonal matrices $C_n > 0$ with positive diagonal entries if

$$C_n(S^n - \mathrm{E}_{\beta^n}[S^n]) \xrightarrow{\mathrm{P}} 0 \text{ for } \beta^n \in \mathcal{B}.$$

The concentration condition requires that there is some normalization (C_n) for which the statistics will concentrate around their means. As we have not scaled statistics

[25]This can be satisfied if there is some geography to nodes and links between distant ones are essentially independent. This is related to the domain increasing asymptotics in the spatial literature, and also to random geometric graphs (RGG, see, e.g., **?**), as well as a variety of latent space models. Without a natural embedding of nodes in some latent space, making assumptions of approximate independence or mixing conditions are stronger than we would like to use here.

[26]Here the subscript notation $\mathrm{E}_\beta[S^n]$ indicates that the expectation takes the probability to be specified by (3.1) with parameters $(S^n, K^n_{S^n}, A_n, \beta)$.

we have to allow for some renormalizations.[27] Viewed as a moment problem, this is essentially just pointwise convergence.

Note that the choice of C_n in the following theorem links them across the two conditions. To guarantee concentration there has to exist a sequence of C_n that goes to 0 fast enough, while to guarantee expectations identification they cannot go to 0 too quickly. So, the C_n's identify the rate at which the statistics approach their means. The key to verifying consistency is then seeing whether there exists such sequences for which both conditions hold simultaneously.

PROPOSITION **1.** *If a sequence of SERGMs* $(S^n, K^n_{S^n}, A_n, \beta^n)$ *is expectations-identified and concentrated with respect to some* C_n*, then the random vectors are consistent so that*

$$|\widehat{\beta}^n(S^n) - \beta^n| \xrightarrow{\mathrm{P}} 0.$$

The proof of Proposition 1 is relatively routine in following standard consistency theorem approaches and appears in the appendix.

It is also useful to state a ratio version of Proposition 1 to address cases where parameters are, for instance, close to 0. Thus, we wish to have a stronger notion of consistency, not only requiring that the estimator approaches the true parameter, but that it does so in terms of a ratio. This requires corresponding definitions of concentration and identification that are ratio based. We do this in Appendix A in Proposition B.1.

The result is fairly tight in that a version of a converse holds as well. In particular, the following result holds.[28]

A sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$, is *rate-expectations-identified* with rates given by a sequence of diagonal matrices $C_n > 0$ with positive diagonal entries if there exist $\gamma_H > \gamma_L > 0$ such that

$$\gamma_H|\beta - \beta^n| > |C_n \mathrm{E}_\beta[S^n] - C_n \mathrm{E}_{\beta^n}[S^n]| > \gamma_L|\beta - \beta^n|$$

for all n.

Rate-expectations-identification is a condition that says that the sequence of diagonal matrices C_n accurately captures the rate at which the expected statistics $\mathrm{E}_\beta[S^n]$ nears $\mathrm{E}_{\beta^n}[S^n]|$ as we let the vector of parameters β approach β^n.

PROPOSITION **2.** *If a sequence of SERGMs* $(S^n, K^n_{S^n}, A_n, \beta^n)$ *is rate-expectations-identified with rates* C^n*, then the random vectors are consistent (i.e.,* $|\widehat{\beta}^n(S^n) - \beta^n| \xrightarrow{\mathrm{P}} 0$*)* if and only if *the sequence is concentrated with respect to* β^n *and the same sequence of* C^n*.*

The above proposition(s) provide high-level conditions wherein the concentration of the sufficient statistics about their expectations, when normalized, delivers consistency. It is then important to know which models will satisfy the concentration requirement.

[27]Notice that the exponential term in the associated likelihood can be written $\exp(S'\beta) = \exp\left(S'C_nC_n^{-1}\beta\right)$ and we are interested in the associated parameters β, not $C_n^{-1}\beta$ which will typically trail off to infinity at polynomial rates in n.

[28]We state a version for Proposition 1, and the analog for ratio-consistency of the type in Proposition B.1 in the appendix is left to the reader.

In the case of iid, time series or even spatial mixing data, we would typically use a concentration inequality such as Chebyshev's inequality or a Chernoff bound to bound the estimate of the mean by some moments and mixing coefficients.

In the appendix we also discuss a characterization of consistency in terms of variance of the statistics in greater detail. The argument is similar to those for standard estimators, e.g. Amemiya (1973): For consistency, we need enough variation so that the system accumulates information and concentrates around its mean. This corresponds to needing the norm of the variance matrix tending to infinity, just as we would in typical regression-like applications.

4.2. **A Class of Consistent and Easily Estimable SERGMs: Simple Count SERGMs.**

To put the above consistency results to work, we now define a natural class of SERGMs that have consistent estimation as well as very straightforward and practical techniques for estimation.

Some ERGMs in which statistics involve counts of subgraphs, can be difficult to estimate, since the associated $N_S(\cdot)$'s can be difficult to calculate. We now discuss an alternative specification, with substituted K_S's that circumvents these difficulties.

Suppose that $S^n = (S_1^n, \ldots, S_k^n)$ is an k-dimensional vector of network statistics whose ℓ-th entry takes on non-negative integer values with a maximum value $\overline{S}_\ell^n \to \infty$. We call such a SERGM specified with $K^n(s) = \prod_\ell \binom{\overline{S}_\ell^n}{s_\ell}$ a *simple count SERGM*. Let let $C_n = \text{Diag}\left\{\overline{S}_\ell^n\right\}_{\ell=1}^k$ be the associated normalizing matrix.

In a simple count SERGM, each statistic can be thought of as counting some aspect of the network: the number of links between nodes of various types, various types of cliques, other subgraphs, the number of pairs of nodes at less than some distance from each other, etc. Note that the class of statistics in a simple count SERGM encompasses most cases of interest. The difference from a standard ERGM is in the specification of the weights K_S that follow a natural counting weight.

Associated with any vector of count statistics S^n on n nodes is a possible range of values. It could be that there are cross restrictions on these values. For example, if we count links S_L^n and isolates S_I^n, then S_L^n cannot exceed $\binom{n-S_I^n}{2}$. In that case

$$A^n = \left\{(s_L, s_I) : s_I \in \{0, 1, \ldots, n\}, s_L \in \left\{0, 1, \ldots, \binom{n - s_I}{2}\right\}\right\}.$$

A sequence of simple count SERGMs (S^n, K^n, A^n, β^n) have statistics that are *not conflicted* if there exists some $\varepsilon > 0$ such that

$$\prod_\ell \left\{\lfloor \mathrm{E}_{\beta_\ell^n}[S_\ell^n](1-\varepsilon)\rfloor, \ldots, \lfloor \mathrm{E}_{\beta^n}[S_\ell^n](1+\varepsilon)\rfloor\right\} \subset A^n$$

for all large enough n.[29]

[29] $\mathrm{E}_{\beta_\ell^n}[S_\ell^n]$ refers to the expectation taken with respect to the one dimensional distribution of S_ℓ^n ignoring the other statistics: i.e., with respect to a SERGM $\frac{K_{S_\ell}^n(s_\ell)\exp(\beta\cdot s_\ell)}{\sum_{s'_\ell \le \overline{S}_\ell^n} K_{S\ell}^n(s'_\ell)\exp\left(\beta\cdot s'_\ell\right)}$. This is slightly different from the condition, since it takes expectations with respect to the unconstrained range of S_ℓ^n rather than cross restrictions imposed under A^n.

Asking that statistics not be conflicted is a fairly simple condition. To see the issue, consider asking for a network on n nodes that has n isolates and a positive number of links is infeasible. Those realizations are conflicted: no network exists with that combination of statistics. The "not conflicted" condition simply asks that at least in some small neighborhood of the unconstrained expected values of the statistics - as if they were each counted completely on their own, they are not conflicted so that they are jointly feasible. So for example, if one expects one third of the nodes to be isolated, and one out of five links to be present, then with large numbers of nodes, these statistics are unlikely to be in conflict. This condition is quite easy to satisfy.[30]

THEOREM **1.** *A sequence of simple count SERGMs that are not conflicted is consistent (i.e.,* $|\widehat{\beta}^n - \beta^n| \xrightarrow{\text{P}} 0$*).*

Moreover, the parameters are asymptotically normally distributed:

$$C_{n,\ell}^{1/2}\left(\widehat{\beta}_\ell^n - \beta_\ell^n\right) \rightsquigarrow \mathcal{N}\left(0, V_\ell^n\right)$$

with $V_{\ell\ell}^n = \left(\frac{\exp\beta}{1+\exp\beta}\right)\left(1 - \frac{\exp\beta}{1+\exp\beta}\right)\frac{\exp\left(\frac{\exp\beta}{1+\exp\beta}\right)}{\left(1+\exp\left(\frac{\exp\beta}{1+\exp\beta}\right)\right)^2}$.

Furthermore, there is a direct estimator available for any model of this type. One can simply estimate parameters p_ℓ from a distribution Bin $\left(\overline{S}_\ell^n, p_\ell\right)$ with s_ℓ as its realization. Then $\widehat{\beta}_\ell := \log\left(\widehat{p}_\ell/(1-\widehat{p}_\ell)\right)$.

It is important to emphasize that this class of SERGMs still allows for strong interdependencies and correlations in link appearances, both within and across statistics. What it takes advantage of is a local approximation of such SERGM distributions in unconflicted regions.

Theorem 1 tells us that unconflicted simple count SERGMs form a consistently estimable class whose statistical properties we understand very well. An advantage of such SERGMs is that they have easy-to-understand and important statistical properties. More generally, this exercise points out that we could also ask about which sorts of other K_S reference distributions lead to consistency of the associated SERGM. We did this above with a very natural class, but one could work with others that embody other ways in which a formation process might lead to weights on statistics.

5. A Class of Estimable Sparse Network Formation Models

We now outline another new model of network formation that is particularly well-suited for estimation in the case of sparse networks. This model also takes its inspiration from the idea that subnetwork configurations are the center of the formation process. Here a network is directly formed through the formation of subnetwork configurations: various cliques, links, stars, etc., form, and the resulting network is their union.

[30]There are other things also embodied in the condition, as there are certain counts of statistics that might not be feasible: for instance it is not possible to have a network with only one triangle missing: once one triangle is removed it also removes many others from the network. Thus, the range of some statistics is not a connected (containing all adjacent entries) subset of the integers. Nonetheless, for lower values of triangles, this is not an issue. Generally, in the relatively sparse ranges of networks that are often of empirical interest, this condition can be easily satisfied.

The reason for considering this class of models is that it is distinct from the simple count SERGMs, in that the network is directly formed from the subgraphs that are counted. Thus, the network is directly formed, rather than having the statistics, and then the network being ancillary. Both are valid sorts of models, and distinct enough that it is valuable to present both.

Moreover, as we show, in sparse settings there are simple and accurate (consistent) estimation techniques for this class of other class of models based on binomial formulas.

These models are widely applicable and estimable, since many social and economic networks of interest are sparse in the sense that the fraction of links that are present versus those absent is small. There are billions of people in the world and typically people have ongoing relationships with at most hundreds of others; there are tens of thousands of researchers in economics and typically authors co-author with handfuls of others; a typical village may have a thousand members each with approximately twenty friends; and so forth.

5.1. **A Subgraph Generation Model: SUGM.**

We now define this other model of network formation, that we sometimes refer to as a "SUGM" or a "Subgraph Model" or "Subgraph Generation Model".

A *subgraph statistic* $S^n_\ell(g)$ is a count of how many instances of some particular (path-connected) subnetwork on m_ℓ nodes appear in g.[31] For example it could be a count of links, triangles, stars comprising m nodes ("$m-1$ stars"), m-cliques. These could also be more intricate subgraphs, for example, groups of m nodes that contain at least one node with characteristics X, another node with characteristics X', etc., and at least some number of links. These could also be directed subgraphs in the case of a directed network.[32]

Consider a vector of subgraph statistics $S^n = (S^n_1, \ldots, S^n_k)$, and corresponding parameters $p^n = (p^n_1, \ldots, p^n_k)$.

A network is randomly formed as follows. First, each of the possible subnetworks counted by S^n_1 is independently formed with a probability p^n_1. Next, each of the possible subnetworks counted by S^n_2 is independently formed with a probability p^n_2. This process continues, through each of the different types of subnetworks counted under S^n. This results in a network g. So, the process is one where various subgraphs are formed, including whatever cliques, stars, links, or other particular network structures that we think could be an important part of a network formation.

So, as an example with triangles and links not in triangles, triangles are formed, and then links that are not already part of any triangle (including incidentally generated triangles) are formed. These could depend on node characteristics: there might be different probabilities for triangles forming based on which types of nodes are in the possible triangle, and similarly for links.

This differs from the SERGM model from Section 4.2 because the actual statistics generated are not directly observed. In particular, there can be a bunch of incidentally generated subnetworks. So the actual counts of statistics under the resulting g can differ from the number that were formed directly under the process. Backing out how

[31]Formally, there is a set $\mathcal{H} = \{H_1(X), ..., H_k(X)\}$ of *representative graphs*, possibly depending on covariates X, of m_ℓ nodes for $\ell = 1, .., k$. S^n_ℓ is given by the number of injective homomorphisms $H_\ell \hookrightarrow g$.

[32]This definition does not admit isolated nodes, but those can also easily be counted.

many were actually formed directly is important in estimating the true parameters of the model, the p_ℓ^n's.

We show that such a model can be easily and consistently estimated if it is sparse enough, in a well-defined way.

5.2. **Incidentally Generated Subgraphs.**

To see why sparsity is important in estimation of SUGMs, consider the following challenge. Suppose that triangles are formed as part of a process of forming a network. In order to estimate the network formation process, we need to know how many triangles were formed. As an example, suppose that the triangles that were directly formed (the cliques that formed) were $(1,2,3)$, $(3,4,5)$, $(1,4,6)$, and $(7,8,9)$. There were really "four" triangles formed; however, we end up with an additional triangle: $(1,3,4)$. Thus it is impossible to include those four triangles without ending up with a fifth triangle. This sort of problem means that some of the arrangements of 4 triangles result in some extras, and so are not really arrangements that have four triangles. This is a challenge in estimating a parameter related to triangle formation, since some of the triangles that we observe were truly generated in the formation process, and others were "incidentally generated".

The key insight in terms of estimation, is that in cases where networks are sparse, then the fraction of incidentally generated subgraphs vanishes.

In order to state our results on this class, we first need a few definitions. Consider a sequence of random graph models indexed by n, each depending on k-dimensional statistics S^n, where k is fixed for the sequence.

For each $\ell \in \{1,\dots,k\}$, let G_ℓ^n be the set of all subnetworks counted by a subgraph statistic S_ℓ^n. In particular, let $g \in G_\ell^n$ be listed as a set of links.[33]

Let us say that a vector of subgraph statistics $S^n = (S_1^n, \dots S_k^n)$ is *nicely-ordered* if the subnetwork counted by S_ℓ^n cannot be a subnetwork of the subnetwork counted by $S_{\ell'}^n$ for $k \geq \ell' > \ell \geq 1$ and $S_{\ell'}^n(g)$ counts subnetworks that are not subnetworks of subnetwork already counted by some $\ell < \ell'$:

(*i*) $g_\ell \in G_\ell^n$ and $g_{\ell'} \in G_{\ell'}^n$ implies that $g_\ell \not\subset g_{\ell'}$

(*ii*) $S_{\ell'}^n(g) = |\{g_{\ell'} \in G_{\ell'}^n : g_{\ell'} \subset g \text{ and } g_{\ell'} \not\subset g_\ell \text{ for any } g_\ell \in G_\ell^n \text{ such that } g_\ell \subset g \text{ for some } \ell < \ell'\}|$.

Note that any set of subgraph statistics can be nicely ordered: simply order the statistics so that $\ell' > \ell$ implies $m_{\ell'} \leq m_\ell$.

So, for example with triangles and links: first we count triangles and second we would count links that are not part of any triangles. The idea is that links are generated through the generation of triangles, and then also through direct generation. We then wish to estimate which things were really generated as triangles and which things were really directly generated as links that are not part of some triangle. On sparse enough graphs, the contamination between these two will be minimal as there will be few things that we mistakenly identify as triangles that were really incidentally generated by other triangles and links. Of course, links that are not part of any triangle could not have been incidentally generated. Thus, the nice ordering tells us that we should first count triangles and next count links that are not part of some triangle. If we were

[33]For example, if g is a triangle, then $G_\ell^n = \{g : g = \{ij, jk, ki\}\}$, where we let ij indicate a link between nodes i and j and exclude self links so that i, j and k are distinct.

counting various cliques, we would first count the largest cliques, then the next largest that are not part of any largest clique, and so forth.

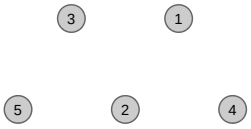
The following definitions allow us to be precise about incidental generation and sparsity.

Consider some subgraph $g_\ell \in G^\ell$. We say that g_ℓ is *incidentally generated* by a set of subgraphs $\{g^j\}_{j\in J}$, indexed by J, where $g^j \in G^{\ell_j}$ for each j for some $\ell_j \geq \ell$,[34] if

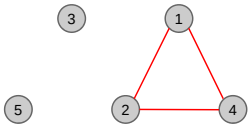
(*i*) $g_\ell \subset \cup_{j\in J} g^j$, and
(*ii*) there is no $j' \in J$ such that $g_\ell \subset \cup_{j\in J, j\neq j'} g^j$

Part (i) states that the subgraph is incidentally generated, and part (ii) of the condition ensures that the set of generating subgraphs is minimal.[35]

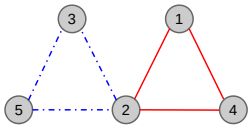
As an example, suppose that we are counting triangles and single links, so that $S_2^n(g)$ is the number of triangles that are a subset of g, and $S_1^n(g)$ is the number of links that are not part of any triangle that are a subset of g. The triangle $\{12, 23, 31\}$ could be generated by the networks g^1, g^2, g^3 where $g^1 = \{12, 24, 41\}$, $g^2 = \{23, 25, 53\}$ and $g^3 = \{31\}$.



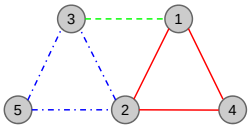
(A) n nodes



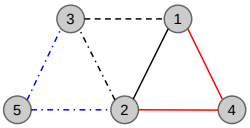
(B) Triangle 124 forms



(C) Triangle 235 forms



(D) Link 13 forms



(E) Triangle 123 is incidentally generated

FIGURE 2. This figure shows the emergence of an incidentally generated triangle.

Figure 2 provides an illustration.

[34]It is possible that $\ell_j = \ell_{j'}$ for some j and j', so that some j's could correspond to different subgraphs of the same type - e.g., a triangle being incidentally generated by other triangles.

[35]Of course, a subgraph could be generated in multiple ways at once. The minimality condition identifies minimal classes that are needed for generation.

5.2.1. *Generating Classes.*

In order to define sparsity, we have to keep track of the various ways in which a subnetwork counted by some statistic in $S^n = (S_1^n, \ldots, S_k^n)$ could have been incidentally generated.

There are typically many ways in which $g_\ell \in G_\ell^n$ could be incidentally generated, but many of them are equivalent up to relabelings. For instance, there are many different combinations of triangles and edges that could incidentally generate a triangle $g_\ell = \{12, 23, 31\}$, however there are only eight ways in which it be done if we ignore the labelings of the nodes outside of g_ℓ: link 12 could be generated either by a triangle or link, and same for link 23, and 31, leading to $2^3 = 8$ ways in which this could happen.

Consider some $g_\ell \in G_\ell^n$ on a set of nodes $N_\ell \subset N$ that is incidentally generated by a set of subnetworks $\{g^j\}_{j\in J}$ and also by another set $\{g^{j'}\}_{j'\in J'}$. We say that $\{g^j\}_{j\in J}$ and $\{g^{j'}\}_{j'\in J'}$ are equivalent generators of g_ℓ if for each g^j there is $g^{j'}$ such that $\ell_j = \ell_{j'}$ and $g_j \cap g_\ell = g_{j'} \cap g_\ell$. So the generating sets play the same roles in g_ℓ but might involve different nodes outside of $N(g_\ell)$. Note that equivalent sets of generators must have the same cardinality as they must both be minimal and involve the same intersections with g_ℓ.

Given this equivalence relation, there are equivalence classes of generating sets of networks for any g_ℓ. It is easy to verify that there are at most $M_\ell \leq k^{m_\ell}$ equivalence classes of (minimal) generating sets for any subnetwork g_ℓ. For each class in the equivalence classes we have some list $(\ell_j, k_j)_{j\in J}$ of the types of subnetworks and the number of nodes that the given subnetwork has intersecting with g_ℓ. We call these the *(minimal) generating classes* of a network g_ℓ and note that these are the same for all members of G_ℓ^n, and so we refer to them as the generating classes of ℓ.

Figure 3 provides an illustration.

5.2.2. *Relative Sparsity.* Given a set of nicely ordered subgraph statistics $S^n = (S_1^n, \ldots, S_k^n)$ and any $\ell \in \{1, \ldots, k\}$ and any generating class $(\ell_j, k_j)_{j\in J}$ of ℓ. Let $M_J = (\sum_{j\in J} k_j) - m_\ell$:[36] For example, in forming a triangle from any combination of triangles and links, each $k_j = 2$ and so $M_J = 6 - 3 = 3$.

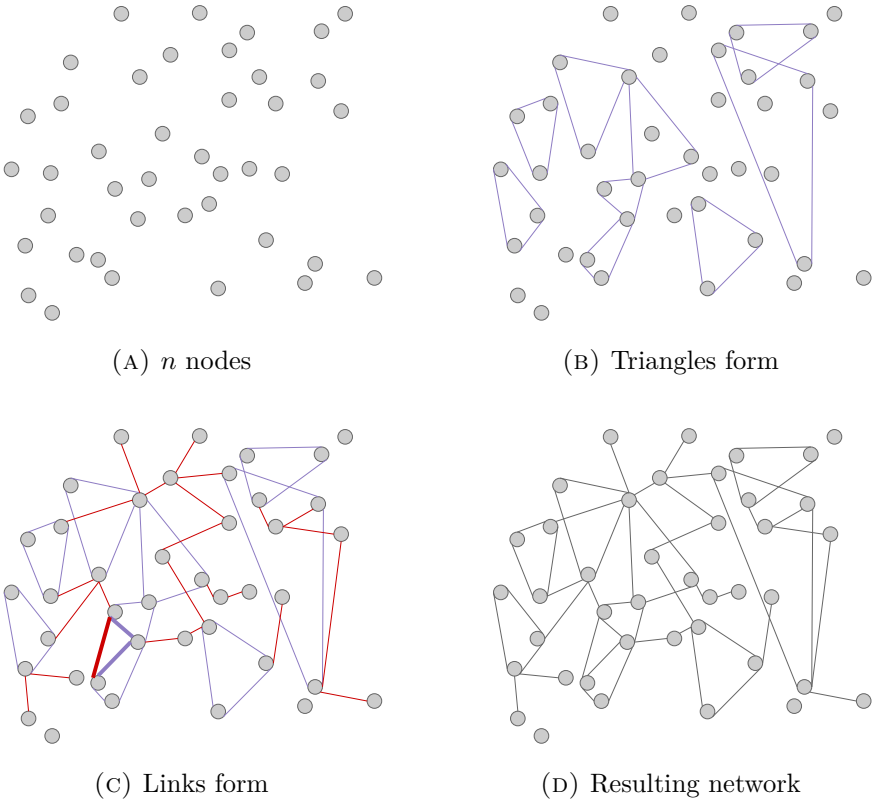
We say that a sequence of models as defined in Section 5.1 with associated nicely-ordered subgraph statistics $S^n = (S_1^n, \ldots, S_k^n)$ and parameters $p^n = (p_1^n, \ldots, p_k^n)$ is *relatively sparse* if for each ℓ and associated generating class $(\ell_j, k_j)_{j\in J}$:

$$\frac{\prod_{j\in J} \mathrm{E}_{p^n}(S_{\ell_j}^n)}{n^{M_J} \mathrm{E}_{p^n}(S_\ell^n)} \to 0.$$

To make this concrete, consider our example with triangles and links. The only generating classes to check in the condition are those of triangles. A triangle can be generated by other combinations of links and triangles. The expected number of triangles that nature generates directly is $\mathrm{E}_{p_T}[S_T^n] = p_T \binom{n}{3}$ and the number of links not in triangles is (approximately) $\mathrm{E}_{p_L}[S_L^n] = p_L \left(\binom{n}{2} - O\left(p_T \binom{n}{3} \right) \right)$. Thus it must be that for each generating class,

$$\frac{\prod_j \mathrm{E}_{p^n}(S_{\ell_j}^n)}{p_T n^6} \to 0.$$

[36]Note that $M_J \geq 1$ since $J \geq 2$ and each set of k_j nodes intersects with at least one other set of $k_{j'}$ nodes for some $j' \neq j$.



(A) n nodes (B) Triangles form

(C) Links form (D) Resulting network

FIGURE 3. The network that is formed on n nodes is shown in panel D. First, triangles form independently with probability p_T^n. Next, links form independently with probability p_L^n on the remaining part of the graph. In panel C we see one incidence of an extra triangle that has been generated by this process.

For the generating class of all triangles, this implies that $p_T^2 n^3 \to 0$, so $p_T = o(n^{-3/2})$. For the generating class of all links, this implies that[37] $p_L^3/p_T \to 0$, which is the obvious condition that triangles formed by independent links are rare compared to triangles formed directly. This implies that (but is not necessarily implied by) $p_L = o(n^{-1/2})$. The remaining generating class combinations are then implied by these ones.

For example, letting $p_T = a(n)/n^2$ and $p_L = b/n$, where $a(n) = o(n^{1/2})$ satisfies the sparsity conditions and leads to an expected degree of $b + a(n)/3$ and an average clustering of roughly $\frac{a(n)}{6(b+a(n)/3)(b+a(n)/3+1)}$. This can be consistent with various clustering rates, and admits rates of links and triangles found various observed networks.[38]

5.3. **Estimation of Sparse Models.**

S^n denotes the actual number of subnetworks of various types that are truly generated, which are not observed by the observer since the resulting g may include incidentally generated subnetworks. Let $\widehat{S}^n(g)$ the observed counts including the incidentally

[37]Given that $p_T = o(n^{-3/2})$, it follows that $\mathrm{E}_{p_L}[S_L^n] = p_L\left(\binom{n}{2} - O\left(p_T\binom{n}{3}\right)\right)$ is proportional to $p_L n^2$.
[38]To match very high clustering rates the model can be altered to include cliques of larger sizes.

generated subnetworks. So, in Figure 3, $S_T^n = 9$ but $\widehat{S}^n(g) = 10$ and from observing g there is no way to know exactly what S_T^n is. Nonetheless, as we prove, under the sparsity condition we can accurately estimate the true statistics, and thus the true parameters under the sparsity conditions.

Before stating our next result, we define the following notation.

Let $\overline{S}_\ell^n(g)$ be the maximum count of S_ℓ^n that is possible on network g. So, for instance if we are counting triangles and links not in triangles, then $\overline{S}_T^n(g) = \binom{n}{3}$ and $\overline{S}_L^n(g) = \binom{n}{2} - L_T(g)$ where $L_T(g)$ is the number of links in triangles in g.[39]

Let

$$\widehat{p}_\ell^n(g) = \frac{S_\ell^n(g)}{\binom{\overline{S}_\ell^n(g)}{S_\ell^n(g)}}$$

and let $D_n = \text{Diag}\left\{p_\ell^n(g) n^{m_\ell}\right\}_{\ell=1}^k$.

So, $\widehat{p}_\ell^n(g)$ is the fraction of possible subgraphs counted by S_ℓ^n that are observed in g, out of all of those that could possible exist in g. This is a direct estimate of the parameter p_ℓ^n, as if these subgraphs were each independently generated and not incidentally generated. D_n is a normalizing matrix.

In order to have $\widehat{p}_\ell^n(g)$ be accurate in the limit two things must be true. First, the network is relatively sparse, which limits the number of incidentally generated subgraphs. And, second, it must be that the potential number of observations of a particular kind of subgraph grows as n grows. This would happen automatically in a sparse network setting if we were simply counting triangles and links not in triangles. However, if nodes have different characteristics (say some demographics), and we are counting triangles and links by node types, then it will also have to be that the number of node with each demographics grows as n grows. If there are never more than 20 nodes with some demographic, then we will never have an accurate estimate of link formation among those nodes.

So, we say that a SUGM model is growing, if the probability that $S_\ell^n(g) \to \infty$ for each ℓ goes to 1.

THEOREM **2.** *Consider a sequence of growing subgraph models as defined in Section 5.1 that are relatively sparse and have associated nicely-ordered subgraph statistics* $S^n = (S_1^n, \ldots, S_k^n)$ *and parameters* $p^n = (p_1^n, \ldots, p_k^n)$. *The estimator is ratio consistent:*[40] $\frac{\widehat{p}_{S_\ell^n}(g)}{p_{S_\ell^n}} \xrightarrow{\text{P}} 1$ *for each* ℓ. *Moreover,*[41]

$$D_n^{1/2}\left((\widehat{p}_1^n, ..., \widehat{p}_k^n)' - (p_1^n, ..., p_k^n)'\right) \rightsquigarrow \mathcal{N}(0, I).$$

Theorem 2 illustrates that nicely-ordered and relatively sparse SUGMs are consistently estimable via easy estimation techniques: ones that are direct and so trivially computable.

[39] In sparse networks, $L_T(g)$ would be vanishing relative to $\binom{n}{2}$ and so could be ignored. And typically, in sparse networks, $\overline{S}_L^n(g)$ will be well approximated by $y_\ell \binom{n}{m_\ell}$, where y_ℓ is the number of possible different subgraphs of type ℓ that can be placed on m_ℓ nodes (e.g., y_ℓ is 1 for a triangle, m for a star on m links, etc.).

[40] This, of course, implies consistency. But given that the parameters might be converging to 0, the ratio version is important.

[41] I is the k-dimensional identity matrix.

The proof of the theorem involves showing that under the nicely-ordered and sparsity conditions the fraction of incidentally generated subnetworks vanishes for each ℓ, and so the observed counts of subnetworks converge to the true ones. Given that these are essentially binomial counts, if the rate at which the incidental counts are generated are limited sufficiently, then as the second part of the theorem states, a variation on a central limit theorem applies, and then normalized errors in parameter estimation are normally distributed, and we know the rates at which the parameters converge to their limits. For inference and tests of significance for single parameter values we note that analytic estimates of the variances are directly computable from the analytic expression of the diagonal of the variance matrix. Of course, more complex inferential procedures and tests can be executed through a simple parametric bootstrap as the model is easily simulated.

To make the convergence rates concrete, take our example with links and triangles, letting $p_T = a/n^2$ and $p_L = b/n$. (This will not be a tight bound, but an example of a realistic level of sparsity that satisfies our conditions for asymptotic normality.) Then one can check the incidental generations for triangles is $o_p(n^{1/2})$, which means that the *fraction* of incidentally generated triangles is $o_p(n^{-1/2})$.

Again, we emphasize that although the estimator here is based on binomial approximations, a SUGM still incorporates interdependencies directly through the subgraphs that are generated. The results make use of the fact that in sparse settings, the picture of interdependencies is clear and are measured by the statistics one-by-one.

6. Simulation Exercises

We numerically illustrate the SUGM approach in several ways. We conduct two exercises using data generated under our model. In both exercises, we use the same class of models: a nicely-ordered sparse SUGM with triangles, links and (possibly) groups.

6.1. **Counting Triangles.**

In the first exercise, we compare estimation of a SUGM to an alternative that does not directly account for triangles, but only captures them indirectly through covariates. That is, the alternative model is one that only accounts for links on a bilateral basis, and triangles form because three individuals were all very close to each other in their characteristics (e.g., location, age, etc.) and so happen to form a triangle.

This exercise illustrates an important bias that emerges if one does not properly account for direct triangle formation: homophily effects are dampened since cross group links are sometimes needed to form appropriate triangles.

This simulation is based on a model in which every node belongs to one of two groups. Due to the presence of homophily, a node is more likely to form a triangle with two other nodes of the same group than to form a triangle where there is heterogeneity in group memberships. Similarly, a node is more likely to form a link with another node of the same group.

Every triple of nodes is designated as either "within" or "across," where a within-triple is such that the three nodes are in the same group, and across is any other triple. Similarly, every pair of nodes is characterized either as "within" or "across".

The nicely-ordered sparse SUGM is as follows. In the first step, triangles are formed independently with probability $p_T^w = 20/n^2$ for within-triples and $p_T^a = 2/n^2$ for across-triples. Next, links are formed with probabilities $p_L^w = 4/n^{1.1}$ and $p_L^a = 1/(4n^{1.1})$ for within and across links, respectively. As such, the sparse SUGM is characterized by $p = (p_L^w, p_L^a, p_T^w, p_T^a)$ (or correspondingly $\beta_k = \log\left(\frac{p_k}{1-\theta_k}\right)$ for a SERGM form).

We report results from two estimation procedures. In the first procedure, we directly estimate $\widehat{p}_T^w$ and $\widehat{p}_T^a$ as the share of observed within and across triangles that have formed, respectively, and then estimate $\widehat{p}_L^w$ and $\widehat{p}_L^a$ as the share of observed remaining within and across links, respectively. Our theoretical results have demonstrated that such an estimation technique is asymptotically consistent. (We also compute corresponding $\widehat{\beta}$'s.)

In our second estimation procedure, we misspecify the model to capture the interdependency through *observed* homophily alone. In this case, link formation as comes from $\theta = (p_L^w, p_L^a)$, alone.

In our simulations we set $n = 500$ and use 100 simulations. Results of the parameter estimates are presented in Table 2.

Model Parameter Estimates

Panel A: Estimation Results for p				
	p1	p2	p3	p4
Truth	0.0043	0.0003	0.0001	0.000008
SubGM	0.004	0.0003	0.0001	7.7907E-06
Homophily Only	0.0173	0.0041	--	--
Absolute Difference Between Estimated to Truth (SubGM)	0.0003	3.0591E-06	0.000024854	2.0932E-07
Absolute Difference Between Estimated to Truth (Homophily)	0.013	0.0038	--	--
Ratio of Estimated to Truth (SubGM)	0.9358	0.9886	0.6893	0.9738
Ratio of Estimated to Truth (Homophily)	4.0218	15.1618	--	--
Panel B: Estimation Results for β				
	β_1	β_2	β_3	β_4
Truth	-5.4455	-8.2221	-9.4334	-11.7361
SubGM	-5.5121	-8.2336	-9.8055	-11.7626
Homophily Only	-4.0406	-5.4995	--	--
Absolute Difference Between Estimated to Truth (SubGM)	0.0666	0.0115	0.3721	0.0265
Absolute Difference Between Estimated to Truth (Homophily)	1.4048	2.7226	--	--
Ratio of Estimated to Truth (SubGM)	1.0122	1.0014	1.0394	1.0023
Ratio of Estimated to Truth (Homophily)	0.742	0.6689	--	--

Notes: Estimation as described in the text, with 100 simulation repetitions performed.

TABLE 2

As expected, we find that the correctly specified SUGM generates accurate results while the misspecified homophily model induces biases. Specifically, the ratio of the within to across linking probabilities is underestimated.

6.2. **Matching Features in Data.** We next examine how well a sparse SUGM with links, triangles, and node characteristics (including geography) does at matching other

features of the network in actual data, when those other features are not directly measured in the model. That is, if the model only uses triangles and links, how well will it reproduce other features of the networks such as component structure, degree distributions, and various eigenvalue properties of adjacency matrices. For our exercises we make use of Banerjee et al. (2012) data-set consisting of networks in 75 Indian villages. Here we focus on two sorts of networks: risk-sharing networks and friendship networks.

In addition to the average degree and clustering (which will be at least partly captured by links and triangles), we are interested in other quantities motivated by theory. We look at the first eigenvalue of the adjacency matrix, which is a measure of diffusiveness of a network under a percolation process (e.g., Bollobás et al., Jackson (2008)). We are also interested in the second eigenvalue of the stochasticized adjacency matrix. This is a quantity that is key in local average learning processes and modulates the time to consensus (Golub and Jackson (2012)). Finally, we look at the fraction of nodes that belong to the giant component of the network, as empirical networks are often not completely connected.

Our procedure is as follows. For every village, for each type of network (risk-sharing or friendship), we estimate each of two network formation models. The first network formation model is a simple sparse SUGM with triangles, links, and geographic covariates (a dummy for being within the median distance in the village). The second network formation model is a simple conditional edge independent link formation model where the probabilities also depend on the geographic covariates. Given our estimated parameters for the village network, from each model generate 100 draws from the SUGM model and the homophily model. We then study the properties of these generated networks across the aforementioned network characteristics: degree, clustering, fraction of nodes in the giant component, maximal eigenvalue, second eigenvalue of the stochasticized matrix.

Table 3 presents the results of our simulation exercise. We show that networks simulated from our estimated parameter values according to the sparse SUGM better display the structural properties exhibited by the empirical Indian village networks than those simulated from a homophily-based model.

Panel A presents results for the risk-sharing networks and Panel B shows results for the friendship networks. Both the sparse SUGM and the homophily models do quite well for average degree. As expected, the homophily model does extremely poorly when it comes to clustering while the SUGM does quite well.

Interestingly, conditioning on the triangles in the SUGM is enough to deliver matches on several other dimensions. Let us focus on the risk-sharing networks for exposition. The fraction of nodes in the giant component is very close between the empirical data and the SUGM (77% to 74%), as is the first eigenvalue (5.23 to 5.26), as well as the second eigenvalue (0.94 to 0.95). Meanwhile, the homophily model performs rather poorly: overestimating the giant component (88%), underestimating the first eigenvalue (3.9) as well as the second eigenvalue (0.28). These patterns hold for the friendship networks as well.

Ultimately, this suggests that there is substantial value added simply by modeling the formation of a sparse number of triangles.

Properties of Networks

Panel A: Risk-sharing networks

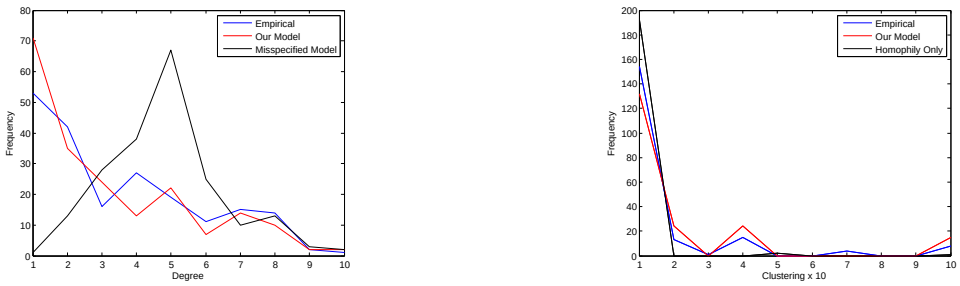
	Truth	SubGM	Homophily Model
Average Degree	2.695	2.524	2.5353
Average Clustering	0.1296	0.1131	0.0107
Fraction in Giant Component	0.7729	0.7403	0.8771
First Eigenvalue	5.2329	5.2555	3.9019
Second Eigenvalue of Stochastized Matrix	0.9422	0.9506	0.2841

Panel B: Friendship networks

	Truth	SubGM	Homophily Model
Average Degree	3.3337	3.4109	3.3864
Average Clustering	0.1073	0.1295	0.0189
Fraction in Giant Component	0.8131	0.8343	0.9599
First Eigenvalue	6.5738	6.3715	4.6249
Second Eigenvalue of Stochastized Matrix	0.9193	0.9153	0.2792

Notes: Estimation as described in the text, with 100 simulation repetitions performed.

TABLE 3



(A) Degree distributions (B) Clustering Coefficients

FIGURE 4. Examples of degree and clustering distributions.

7. EXTENSIONS AND APPLICATIONS

We now describe a few special cases of SERGM and SUGM models that illustrate their use in analyzing strategic network formation. We also apply the model to data.

7.1. **Strategic Network Formation.** As mentioned in Section 2.6, an important use of SERGMs (and SUGMs) is for estimating strategic network formation models. While a full treatise on this subject would take us away from the focus of this paper, we illustrate how the models can be used in some contexts, with extensions to other contexts left to the reader and further research.

7.1.1. *Mutual Consent Formation Models.* Here we describe a strategic network model that harnesses some of the power of Theorem 2. The key aspect of the model is

that decisions to link are not only bilateral but instead multilateral: sub-groups of individuals decide whether or not to form cliques. A pair of individuals may meet and decide (mutually) as to whether to add a single link, but also a group of three (or more) may meet and decide whether to form a clique. Moreover, the probability that they form the clique could depend upon the characteristics of the individuals involved.

Consider cliques $CL_m \subset \{1, \ldots, n\}$ of m nodes with characteristics $X = (X_i)_{i \in Cl_m}$ meeting at random. There are certain aspects, $h_i(X)$, that affect i's benefits from the clique.[42]

There is a probability $\pi_m(X)$ that a clique of m individuals with characteristics X meets. So it might be more likely that individuals of similar ages meet than ones with different ages.

Individual i obtains a utility

$$U_{i,Cl_m} = \gamma_{m,X_i} h_i(X) - \epsilon_{i,Cl_m}$$

from the formation of a given clique Cl_m, where γ_{m,X_i} depends on the size of the clique and possibly on the characteristics of i and ϵ_{i,Cl_m} is a random idiosyncratic term.

The clique then forms conditional upon it having met if and only if $U_{i,Cl_m} \geq 0$ for all i (say with at least one strictly positive).[43] If the error term has an atomless distribution, then the strictness is inconsequential.

Let $F_m(\cdot)$ be the distribution of error terms when the clique is of size m.

So, the probability that a clique Cl_m of size m with characteristics $X = (X_i)_{i \in Cl_m}$ is formed is

$$p_{m,X} = \pi_m(X) \prod_{i \in Cl_m} F_m \left(\gamma_{m,X_i} h_i(X) \right).$$

Suppose that the largest cliques that can form are of some size M. Let $S_{m,X}(g)$ denote the number of cliques of size $m \in \{2, \ldots, M\}$ consisting of individuals with characteristics X that form in a network g, counted excluding cliques that are already a subset of some $m' > m$.[44] Under the conditions of Theorem 2, the SUGM in (7.1.1) is then easily estimated and consistent.[45]

[42] For instance if X were a list of the individuals' ages, then it might be that i's benefit from the clique is a function of i's distance from the average characteristics: $h_i(X) = \left\| X_i - \sum_{j \neq i} X_j/(m-1) \right\|$. It could also be that i benefits from the maximum value of X_{-i}, or suffers from variation in the characteristics. h can be tailored to the specific application and list characteristics.

[43] This then corresponds nests pairwise stability as defined by Jackson and Wolinsky (1996), subject to the meeting process. One can adjust this to take into account other rules for group formation.

[44] For instance if there is a group of four identical individuals in a clique, then we will not also count cliques of size three that are a subset of that, but only count cliques that are not already subset of some larger clique.

[45] Although the probabilities of various cliques are directly estimable (and hence identified) under the conditions of Theorem 2, of course whether the various parts of $\pi_m(X) \prod_{i \in Cl_m} F_m \left(\gamma_m h_i(X) \right)$ are well identified depends on the specifics of the m and the functional forms involved. Just as an example, consider a situation with two types. Set $X_i = 1$ if i is of type 1 and $X_i = 2$ if i is of type 2 and $h_i(X) = X_i - X_j$ and consider $m = 2$. Then γ_2 and $-\gamma_2$ both lead to the same value of $F_2 \left(\gamma_2 (X_i - X_j) \right) \times F_2 \left(\gamma_2 (X_j - X_i) \right)$. So here it could not be judged whether the type 2 has a greater expected utility (net of the random term) from the match than the type 1 or whether it is the reverse. There are some obvious simplified formulations that allow for identification, for example setting instead $h_i(X) = |X_i - X_j|$. It might also require specifying a (nonlinear) functional form for π_m as in Currarini et al. (2010).

Note that the above formulation centers on cliques, but substituting H_m to be some particular subgraph instead of the clique, and keeping track of each node's position in the subgraph.

It is important to note that such a formulation allows us to do welfare computations and changes in welfare due to changes in, say, the distribution of X if they include parameters that a policymaker may control - or it may be that a policymaker could change the π function in some well-defined way by say, subsidizing interactions among groups with certain sorts of characteristics who might tend to meet infrequently.

7.1.2. *Search Intensity Models.* Although the above model allows for strategic formation, it is a model where there is no direct strategic interaction between agent's decisions to form subnetworks. Agents effectively have dominant strategies to try to form all of the cliques that benefit them, and the ones that form depend on the actions of others.

To account more directly for externalities that are present in network formation processes, we can also include search intensities as have been analyzed in various formation models such as Currarini et al. (2009), Currarini et al. (2010), and Golub and Livne (2010). Those models are of bilateral (independent) link formation; but are easily extended more general SUGMs as we briefly describe.

Each agent i with characteristics X_i puts in a search effort $e(X_i, m, X) \in [0, 1]$ to form cliques of size m with characteristics X. $m = 2$ indicates links, and so $e(X_i, 2, (X_i, X_j))$ is the effort that agent i expends in trying to form links with agents who have characteristics X_j; and $e(X_i, 3, (X_i, X_j, X_h))$ is the effort that agent i expends in trying to form triangles where the other two agents have characteristics X_j and X_h, and so forth.

"Effort" is simply a shorthand for either the time spent socializing in various ways, or else it could simply indicate a relative openness to forming relationships of various types.

An agent obtains a utility

$$u(X_i, m, X)$$

from being of type X_i in a clique of size m with characteristics X.

The probability that a given clique Cl_m forms depends on the vector of efforts for such cliques of those in the clique, $(e_j(, m, X))_{j \in Cl_m}$, according to a function $\pi_{m,X}((e_j(m, X))_{j \in Cl_m})$ that is nondecreasing in each of its arguments.

An agent also pays a cost of network formation: $c(X_i, (e_i(m, X))_{m,X})$ that depends on his or her characteristics X_i and the search efforts that he or she exerts in forming various links and cliques, $(e_i(, m, X))_{m,X}$.

Thus, an agent i's overall payoff as a function of the all of the agents' efforts is described by

$$\left(\sum_{m, Cl_m : i \in Cl_m} \pi_{m,X}((e_j(m, X))_{j \in Cl_m}) u(X_i, m, X) \right) - c(X_i, (e_i(m, X))_{m,X})$$

In a case where the u's are nonnegative, this defines a supermodular game: agent i's change in payoff from increasing any dimension of $(e_i(m, X))_{m,X}$ is nondecreasing in the vector of strategies $(e_j(m, X))_{j \neq i, m, X}$. In such games, pure strategy equilibria exist and form a complete lattice (e.g., see Topkis (2001)). Additional conditions on π, u,

and c can ensure uniqueness of equilibrium, depending on the specific functional forms that are used to parameterize the model, or one can appeal to equilibrium selection.[46]

Models of this structure thus define SUGMs, where the relative frequencies $p_{m,X}$ of cliques of size m consisting of agents with characteristics described by the profile X. Specifying functional forms for π, u and c then allows for estimation of parameters of the model and of the equilibrium, provided the specification is tight enough to be well-identified.

Again, although the above formulation is described for cliques, it is easily adjusted for any subgraphs (for instance an agent may value being at the center of a star with m agents, etc.). In the obvious extension one needs to keep track of the positions of the various types of agents in the subgraph as there are then asymmetries in positions and, for instance, agents might care about the characteristics of the agent at the center of a star.

7.1.3. *Other Formulations.*

The above formulations allow for various cliques to form, and thus nontrivial clustering in networks, as well as high frequencies of friends in common even in sparse networks. They are also extend directly to other sorts of subgraphs: stars, lines and other configurations.

Another feature of interest is that an agent may care about friends of friends: and may care to link to others who have more neighbors. For example, that is a central feature in the connections model of Jackson and Wolinsky (1996), and will be a feature of settings where agents care substantially about indirect communication and favors, etc.

Such a formulation can be viewed as an ERGM, as shown by Mele (2011). Specifically, in his model in every period a pair of individuals is called to meet and then the agents decide whether or not they want to link. Their utility from forming the link may depend on the network that has formed through that stage. Mele shows that in a model where agents value links and the indirect connections that they convey in an additively separable form, the invariant distribution of such a process results in an ERGM.[47,48] In particular, his model specification leads to a setting where the relevant ERGM depends on links and two-stars (so friends of friends). By applying our results, such a model can then be estimated directly and easily as a (sparse) SUGM.

Another novel approach to estimating strategic network formation models is to work with ones that can be projected onto subnetworks, as by Sheng (2012). Building on techniques from estimating bounds from games with multiple equilibria as in Ciliberto and Tamer (2009), Sheng derives bounds on what sorts of subnetworks should be observed under pairwise stability. Ultimately, her approach could be used to produce a SUGM which then could be estimated using our techniques.

[46]Here there are positive spillovers/externalities from strategies, and so generally the maximal equilibrium will Pareto dominate the others, and so a standard refinement would be to look at the Pareto efficient equilibrium which is then unique and pure (e.g., see Vives (2007) for some background).

[47]There must be *some* ERGM representation of the invariant distribution by the Hammersley-Clifford theorem. The Mele result shows that the potential function associated with the game essentially defines the corresponding ERGM.

[48]For an alternative specification in terms of a Bayesian equilibrium, see Leung (2012).

7.2. **Noise and Almost-Cliques.**

Although observed social networks often exhibit significant clustering and numbers of triangles significantly above what would be observed with independent link formation, typically full cliques of larger sizes are rarer. There can be different reasons for this, and here we briefly discuss a correction for this.

One reason for a failure to see completed cliques of higher level is that small amounts of measure error makes them exponentially (to the square of the size of the clique) less likely to be observed. To see this point, consider the simplest situation such that there is a probability $\varepsilon > 0$ that a data set fails to exhibit any given link (independently across links). A clique of ℓ nodes has $\ell(\ell-1)/2$ possible links. If it were present, then it would be fully observed only with a probability that all of its links are observed, or $(1-\varepsilon)^{\ell(\ell-1)/2}$. For example, suppose that $\varepsilon = .1$. Then the probability that a clique of size 3 is observed is .73, while this drops to .53 for a clique of size 4, to .35 for a clique of size 5, and to .21 for a clique of size 6.

Why is this an important issue? Suppose that the true model generates cliques of size 3 and 4 in addition to links. The 27 percent of triangles that are missed due to measurement error, will end up classified as two or fewer links, thus biasing downwards the parameter on triangles and upwards that on links. The 47 percent of cliques of size 4 that are missed, will generally contribute to increased observations of triangles and links. For instance, deleting one link from a clique of size 4 makes it appear as two cliques of size three. Taking two links out makes it either lead to a triangle plus a link or four extra links. This can substantially bias upwards the counts of triangles and links.

There are two ways to deal with this, one more precise and one easier but less precise. The precise way to do this, is to model the error and include it in the specification of a SERGM or SUGM. For example, given a SUGM with links, triangles and cliques of size 4 together with probability ε of missing a link in the data, then the probability of seeing a clique of size 4 becomes $p_4(1-\varepsilon)^6$. The probability of seeing a clique of size 4 less one link (under the sparsity conditions) becomes $p_4 6(1-\varepsilon)^5\varepsilon$, and the probability of seeing a clique of size 4 less two links (again under the sparsity conditions) becomes $p_4 15(1-\varepsilon)^4\varepsilon^2$, and so forth. Thus, one can then consider the following count statistics: links, two-stars, triangles, cliques of size 4 less two links, cliques of size 4 less one link, and cliques of size 4, etc. Each of these has a well defined probability given the original specified p_2, p_3, p_4 and the ε. One can then estimate these parameters (including the ε) using the SUGM techniques from our theorem.

A less precise way to do correct for incomplete cliques would be to simply count any structure on ℓ nodes that has at least $x_\ell \ell(\ell-1)/2$ links as a clique of size ℓ, where $x_\ell \leq 1$ is some simple adjustment factor. While a crude adjustment, this could still improve over ignoring the issue altogether. For example, in a setting with links, triangles and 4-cliques, one could make a simple adjustment by counting any configuration on four nodes with at least 4 links as a 4-clique, but then stick with counting triangles only if all three links are present and they are not part of any modified 4 clique.

While measurement error is one possible way in which a clique could form with missing links, another possibility is that the formation process is one where agents choose which relationships to form, but have payoffs to cliques that allow them to benefit

from fewer than all links being present, and allowing some noise in the formation (either utility-based or trembles) that lead to some links in cliques not forming. While different in interpretation, the techniques that account for such missing links would be similar to that described above: simply directly accounting for the probabilities that various subgraphs form. As SERGMs and SUGMs can already admit arbitrary subgraphs as statistics, this would simply place additional restrictions on the relationships between the probabilities of various subgraphs, but would otherwise use similar estimation techniques as outlined in our results.

7.3. **The Use of SERGMs and SUGMs in Policy Analysis.** One motivation for modeling network formation is so that we can also understand what might lead to changes in network structures. This is particularly important in situations where network features are drivers of economic behaviors: for instance, expansion properties of a network can impact diffusion of innovations (e.g., it might be the case that in networks with a higher maximal eigenvalue farmers are significantly more likely to have learned how to properly use a new farming technology that has been introduced). By estimating a network formation model (a SERGM or SUGM), we can understand what sorts of interventions could generate variation in network features of interest.[49] We might find a certain sort of intervention (e.g., something to enhance the meeting process, or programs to encourage relationships across groups) would significantly affect homophily, expansion properties of the network, or some other feature of interest. This can then lead to changes in diffusion of information or other behaviors.

7.4. **Multigraphs and motives for linking.** An important question in analyzing economic and social networks is what sorts of interactions exist between different sorts of relationships: does the fact that two agents borrow and lend money to each other make it more likely that they will also seek advice from each other, or engage in other sorts of relationships? Is it more likely that pairs of agents who have multiple relationships with each other will be embedded in cliques?

The SERGM and SUGM models that we described extend directly to multigraph settings in obvious ways, as we illustrate in the application below.

7.4.1. *Empirical Application: Motives for linking.* The observation that different relationships that individuals may have with each other may be interdependent dates back to Simmel (1918). Here, we outline three basic theories for multiplexing and then investigate them in some data using our results.

First, there may be a fixed cost of making links. Conditional on having established a relationship with another person, it is cheaper to construct the second relationship. That is, if i is willing to link to j to borrow rice, then it might be cheaper to for i to also establish a social link with j.

Second, decisions of individuals to form relationships depend on their characteristics. To the extent that some of those characteristics are not observed in a data set, their decisions to form relationships could be correlated even after adjusting for all observable characteristics.

Third, is what we call a *support theory* based on favor exchange. In such a theory there are two ways to support the exchange of favors. One is to have frequent enough

[49]Thanks to Esther Duflo for emphasizing the importance of this use of the models.

favor exchange so that it is in each individual's interest to provide favors whenever his or her friend asks for one. In this way, an isolated pair of individuals who have multiple relationships (exchanging various types of favors, etc.) can sustain favor exchange because the multiple relationships make the pairing valuable enough to continue exchanging favors. Another way to provide incentives to exchange favors is to embed that exchange in a clique. If some individual in a clique fails to provide a favor when asked, then all of the other individuals in the clique collectively punish the agent by not providing favors in exchange. The embedding in a clique provides the additional incentives necessary to motivate favor provision. Thus, the leverage needed to incentivize the provision of favors comes either through building valuable multi-level bilateral relationships, or else by embedding the exchange within cliques with more individuals. It should be rarer to see isolated favor exchange between two individuals who do not share multiple relationships and do not have friends in common.[50]

To examine these theories, we use the framework outlined in Section 5.1. We use the data collected across 75 Indian villages collected by Banerjee et al. (2012).[51] The specific usefulness of the data is that it includes multigraphs including various forms of favor exchange links (whether households borrow or lend kerosene, rice, money,...) and social links (whether household members visit each other socially, go to temple together, ...). For the purposes of our analysis, we use all 75 Indian villages and construct networks at the household level. We build undirected, unweighted multigraphs with two link types (where a link of a given type is present if either household claimed that type of relationship with the other). Favor exchange links indicate whether households borrow/lend kerosene, rice and other material goods from/to each other and social links indicate whether members of a household visit members of another household or receive these members as guests.

To distinguish between the first two theories, we examine multiplexing of supported links (those in cliques) to those that are not supported. Under a theory of fixed costs of linking, the multiplexing would have no reason to vary, while under a theory of unobserved characteristics they would. In particular, unobserved characteristics should result in triangle formation among trios of nodes that have similar characteristics and then should also have high multiplexing. Unsupported links are more likely to have arisen spuriously and are thus more likely to be unsupported. The evidence is consistent with the similar characteristics theory but not with fixed costs, as more than 70 percent of supported links are multiplex, while less than 20 percent of unsupported links are multiplex, as described in Table 4.[52]

[50]For more on this, see Jackson et al. (2012).

[51]For more background on the data see Banerjee et al. (2012), and for more on favor exchange in the data see Jackson et al. (2012).

[52]The entries of Table 4 are constructed as follows. In columns 1-2 we consider either the favor graph or the social graph and construct the share of links in a village that are supported (members of a triangle). We then compute the mean and standard error across the 75 independent villages. In columns 3-4 we consider either the favor graph or the social graph and construct the share of *unsupported* links that are multiplexed within the village. Again we present the mean and standard error across the 75 villages. In columns 7-8 we consider the share of total links of a given type that are either supported or multiplexed. Again we present the mean and standard error across the 75 villages. Columns 3, 6, and 9 present differences between the favor graph and social graph shares with bootstrapped standard errors.

Support and Multiplexing

Supported/Total			*Multiplexed/Total*			*Supported or Multiplexed / Total*		
Favor	Social	Difference	Favor	Social	Difference	Favor	Social	Difference
0.7303	0.6928	0.0376	0.1907	0.1400	0.0508	0.9210	0.8327	0.0883
(0.0113)	(0.0111)	(0.0026)	(0.0085)	(0.0066)	(0.0028)	(0.0044)	(0.0062)	(0.0042)

Notes: Favor links are borrowing/lending kerosene or rice. Social links are given by those who visit one's house or those whose house the individual visits. In columns 1-2 we look at the share of a given link type that are supported (as well as the difference in column 3). In columns 4-5 we look at the share of a given link type that are unsupported but are multiplexed (as well as the difference in column 6). In columns 7-8 we look at the share of a given link type that are either supported or multiplexed (as well as the difference in column 9). Standard errors computed at the cross network level are presented in parantheses.

TABLE 4

Next, let us examine the support theory. Under that theory, favor exchange links should be more likely to either be embedded in triangles or multiplexed than social links that might not involve explicit favor exchange. Table 4 provides figures for both links that involve explicit favor exchange and those that are more social in nature. It shows that 73 percent of favor exchange links are supported (embedded in triangles) as compared to 69 percent of social links. While they are statistically significantly different, the 3.8 percent gap is small relative to the mean. Next we look at multiplexing of unsupported links. This exercise of counting unsupported links separately is legitimized under our SUGM. We find that while nearly 19 percent of unsupported favor exchange links are multiplexed, only about 14 percent of social links exhibit multiplexing. This 5 percent difference is large relative to the mean. Putting this together, favor exchange links are significantly more likely to be reinforced in some way – either through multiplexing or support – as is consistent with the support theory.

The analysis finds statistics that are consistent with both the similarity in characteristics theory and the support theory, and not with a simple fixed cost theory.[53] Finally, we interpret the fact that favor exchange links are differentially more likely to be multiplexed than social links to be suggestive evidence in favor of the support theory. The identifying assumption here is that any reasonable story of correlated characteristics driving such patterns which are different from favor exchange ought to be symmetric in their effect on multiplexing probabilities, whereas support theory suggests that only unsupported favor links ought to be more likely to be multiplexed.

8. CONCLUSION

We presented two new classes of models of network formation, SERGMs and SUGMs, based on the idea that network formation is driven by the properties of a network or by the formation of various subgraphs. This turns the focus away from the network as the unit of analysis, and instead focuses on its properties, subgraphs and statistics. This perspective allowed us to develop simple and direct estimation techniques and to derive results concerning consistency and asymptotic distributions of the estimators. Given the growing literature on estimation of network formation, such results are essential.

[53]It could be that there are more complex fixed costs that depend on subgraph configurations, but costs for second relationships would have to be significantly lower in cliques than in isolation.

The approach we have taken, can certainly be extended to other formulations. We have discussed the incorporation of several classes of models that allow for random meetings, search, and strategic elements of network formation. Given the numerous theories underlying social structure and network formation, more such models will certainly be needed and the basic methodology developed here should be helpful along that road.

References

Banerjee, A., A. Chandrasekhar, E. Duflo, and M. Jackson (2012): "Diffusion of Microfinance," *NBER Working Paper 17743 http://www.stanford.edu/ jacksonm/diffusionofmf.pdf*.

Bhamidi, S., G. Bresler, and A. Sly (2008): "Mixing time of exponential random graphs," *Arxiv preprint arXiv:0812.2265*.

Bickel, P., A. Chen, and E. Levina (2011): "The method of moments and degree distributions for network models," *Annals of Statistics*, 39, 2280–2301.

Bollobás, B., S. Janson, and O. Riordan (2011): "Sparse random graphs with clustering," *Random Structures & Algorithms*, 38, 269–323.

Bramoullé, Y. and R. Kranton (2007): "Risk-sharing networks," *Journal of Economic Behavior & Organization*, 64, 275–294.

Burt, R. (1992): *Structural Holes: The Structure of Competition*, Harvard University Press.

Calvo-Armengol, A. (2004): "Job contact networks," *Journal of Economic Theory*, 115, 191–206.

Chatterjee, S. and P. Diaconis (2011): "Estimating and Understanding Exponential Random Graph Models," *Arxiv preprint arXiv:1102.2650*.

Chatterjee, S., P. Diaconis, and A. Sly (2010): "Random graphs with a given degree sequence," *Arxiv preprint arXiv:1005.1136*.

Christakis, N., J. Fowler, G. Imbens, and K. Kalyanaraman (2010): "An Empirical Model for Strategic Network Formation," *NBER Working Paper*.

Ciliberto, F. and E. Tamer (2009): "Market structure and multiple equilibria in airline markets," *Econometrica*, 77, 1791–1828.

Coleman, J. (1988): "Social Capital in the Creation of Human Capital," *American Journal of Sociology*, 94, S95–S120.

Currarini, S., M. Jackson, and P. Pin (2009): "An economic model of friendship: Homophily, minorities, and segregation," *Econometrica*, 77, 1003–1045.

——— (2010): "Identifying the roles of race-based choice and chance in high school friendship network formation," *Proceedings of the National Academy of Sciences*, 107, 4857–4861.

Frank, O. and D. Strauss (1986): "Markov graphs," *Journal of the American Statistical Association*, 832–842.

Furusawa and H. Konishi (2007): "Free trade networks," *Journal of International Economics*, 7, 310–335.

Golub, B. and M. Jackson (2012): "How Homophily Affects the Speed of Learning and Best-Response Dynamics," *Quarterly Journal of Economics*, 127, 1287–1338.

Golub, B. and Y. Livne (2010): "Strategic Random Networks: Why Social Networking Technology Matters," *SSRN working paper*.

Hammersley, J. and P. Clifford (1971): "Markov fields on finite graphs and lattices," .

Handcock, M. (2003): "Assessing degeneracy in statistical models of social networks," .

Jackson, M. (2008): *Social and economic networks*, Princeton: Princeton University Press.

——— (2011): *Handbook of Social Economics*, San Diego: North Holland, chap. An Overview of Social Networks and Economic Applications.

Jackson, M., T. Barraquer, and X. Tan (2012): "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102, 1857–1897.

Jackson, M. and A. Wolinsky (1996): "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71, 44–74.

Krackhardt, D. (1988): "Predicting with networks: nonparametric multiple regression analysis of dyadic data," *Social Networks*, 10, 359–381.

Leung, M. (2012): "An Econometric Model of Network Formation with Low Computational Burden," *working paper*.

Mele, A. (2011): "A Structural Model of Segregation in Social Networks," *working paper*.

Moreno, J. and H. Jennings (1938): "Statistics of Social Configurations," *Sociometry*, 1, 342–374.

Rinaldo, A., S. Petrovic, and S. Fienberg (2011): "Maximum likelihood estimation in network models," *arXiv preprint arXiv:1105.6145*.

Schutz, B. (2003): *Gravity from the ground up*, Cambridge Univ Pr.

Shalizi, C. and A. Rinaldo (2012): "Consistency under Sampling of Exponential Random Graph Models," *arXiv preprint arXiv:1111.3054v3*.

Sheng, S. (2012): "Identification and Estimation of Network Formation Games," *USC mimeo*.

Snijders, T. (2002): "Markov Chain Monte Carlo Estimation of Exponential Random Graph Models," *Journal of Social Structure*, 3, 240.

Snijders, T., P. Pattison, G. Robins, and M. Handcock (2006): "New specifications for exponential random graph models," *Sociological Methodology*, 36, 99–153.

Topkis, D. (2001): *Supermodularity and Complementarity*, Princeton University Press.

Vives, X. (2007): "Supermodularity and Supermodular Games," *IESE Occasional Paper 07/18*.

Wasserman, S. and P. Pattison (1996): "Logit models and logistic regressions for social networks: I. An introduction to Markov graphs andp," *Psychometrika*, 61, 401–425.

Appendix A. Proofs

Proof of Proposition 1. Recall that the MLE $\widehat{\beta}^n(s)$ is the β that solves

$$s = \mathrm{E}_\beta[S^n].$$

Thus, since concentration implies that

$$C^n(S^n - \mathrm{E}_{\beta^n}[S^n]) \xrightarrow{\mathrm{P}} 0,$$

it follows that

$$C^n(\mathrm{E}_{\widehat{\beta}^n(S^n)}[S^n] - \mathrm{E}_{\beta^n}[S^n]) \xrightarrow{\mathrm{P}} 0.$$

Given that expectations identification implies that

$$|C^n(\mathrm{E}_{\beta}[S^n] - \mathrm{E}_{\beta^n}[S^n])| > \gamma|\beta - \beta^n|$$

for all n, it follows that

$$|\widehat{\beta}^n(S^n) - \beta^n| \xrightarrow{\mathrm{P}} 0$$

as claimed. ■

These results can be rephrased in terms of standard properties of extremum estimators and identifiable uniqueness.[54]

PROPOSITION **A.1.** *If a sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$ satisfies concentration, rate-expectations identification, and the β^n all lie in a compact set, then $|\widehat{\beta}^n(S^n) - \beta^n| \xrightarrow{\mathrm{P}} 0$.*

Proof of Proposition A.1. Let

$$m_n(\beta) := C_n(S^n - \mathrm{E}_{\beta}[S]).$$

The objective function is $Q_n(\beta) := m_n(\beta)' m_n(\beta)$.

First, we want to show that the moment function satisfies a uniform law of large numbers: $\sup_{\beta \in \mathcal{B}} \|m_n(\beta)\| = o_p(1)$. By concentration, we have pointwise convergence of $m_n(\beta)$ to zero in probability. Therefore, we need to only check stochastic equicontinuity: that for every $\eta > 0$ there is a $\delta > 0$ with

$$\mathrm{P}\left(\sup_{\|\beta - \beta'\| < \delta} |m_n(\beta) - m_n(\beta')| > \eta\right) < \eta.$$

A sufficient condition (Andrews, 1994) is if a Hölder condition is satisfied:

$$|m_n(\beta) - m_n(\beta')| \leq X_n \cdot \|\beta - \beta'\|$$

where X_n is some $O_p(1)$ random variable. This is directly guaranteed by rate-expectations identification with $X_n = \gamma_H$.

Second, one can check that expectations identifiability guarantees identifiable uniqueness. Together with compactness of $\mathcal{B}$, this implies the above implies that $\widehat{\beta}$ is consistent for β^n. ■

Proof of Proposition 2. Recall that the MLE $\widehat{\beta}^n(s)$ is the β that solves

$$s = \mathrm{E}_{\beta}[S^n].$$

Given Proposition 1, we need only show that if consistency holds then concentration must also hold.

[54]Following, e.g., Gallant and White (1988), we say the sequence β^n is identifiably unique on $\mathcal{B}$ if

$$\liminf_{n \to \infty} \inf_{\beta \in B_{\beta^n}(\epsilon)} Q_{0,n}(\beta^n) - Q_{0,n}(\beta) > 0.$$

Given that rate-expectations identification implies that

$$a^n|\mathrm{E}_\beta[S^n] - \mathrm{E}_{\beta^n}[S^n]| < \gamma_H|\beta - \beta_0|$$

for all n, it follows that if if consistency holds so that

$$|\widehat{\beta}^n(S^n) - \beta^n| \xrightarrow{\mathrm{P}} 0,$$

then it must be that

$$a^n|\mathrm{E}_{\widehat{\beta}^n}[S^n] - \mathrm{E}_{\beta^n}[S^n]| \xrightarrow{\mathrm{P}} 0.$$

This implies that

$$a^n|S^n - \mathrm{E}_{\beta^n}[S^n]| \xrightarrow{\mathrm{P}} 0,$$

which implies concentration. ■

Proof of Theorem 1.

Consider a sequence of simple count statistics $S^n = (S_1^n, \ldots, S_k^n)$ whose ℓ-th entry takes on nonnegative integer values with some maximum value $\overline{S}_\ell^n \to \infty$, and the simple count SERGMs specified with $K^n(s) = \prod_\ell \binom{\overline{S}_\ell^n}{s_\ell}$.

$$\mathrm{P}_\beta(S^n = s) = \frac{\prod_\ell \binom{\overline{S}_\ell^n}{s_\ell} \exp(\beta^n \cdot s)}{\sum_{s' \in A^n} \prod_\ell \binom{\overline{S}_\ell^n}{s'_\ell} \exp(\beta^n \cdot s')}. \tag{A.1}$$

We rewrite (A.1) as

$$\mathrm{P}_\beta(S^n = s) = \frac{\prod_\ell \left[\binom{\overline{S}_\ell^n}{s_\ell} \exp(\beta_\ell^n s_\ell)\right]}{\sum_{s' \in A^n} \prod_\ell \left[\binom{\overline{S}_\ell^n}{s'_\ell} \exp(\beta_\ell^n s'_\ell)\right]}.$$

If we consider the distribution conditional on S^n lying in

$$B^n = \prod_\ell \{\lfloor E_{\beta^n}[S_\ell^n](1-\varepsilon)\rfloor, \ldots, \lfloor E_{\beta^n}[S_\ell^n](1+\varepsilon)\rfloor\} \subset A^n$$

(for large enough n under the not conflicted condition), then we can write

$$\mathrm{P}_\beta(S^n = s | s \in B^n) = \frac{\prod_\ell \left[\binom{\overline{S}_\ell^n}{s_\ell} \exp(\beta_\ell^n s_\ell)\right]}{\prod_\ell \left[\sum_{s'_\ell \in B_\ell^n} \binom{\overline{S}_\ell^n}{s'_\ell} \exp(\beta_\ell^n s'_\ell)\right]}$$

or

$$\mathrm{P}_\beta(S^n = s | s \in B^n) = \prod_\ell \frac{\binom{\overline{S}_\ell^n}{s_\ell} \exp(\beta_\ell^n s_\ell)}{\sum_{s'_\ell \in B_\ell^n} \binom{\overline{S}_\ell^n}{s'_\ell} \exp(\beta_\ell^n s'_\ell)}. \tag{A.2}$$

Consider a binomial distribution that has $p_\ell^n = \frac{E_{\beta^n_\ell}[S_\ell^n])}{\overline{S}_\ell^n}$, with range from 0 to $\overline{S}_\ell^n$. The corresponding probability that $\widetilde{S}_\ell^n = s_\ell$ given $s_\ell \in B_\ell^n$ can be written as

$$\frac{\binom{\overline{S}_\ell^n}{s_\ell} \exp(\beta_\ell^n s_\ell)}{\sum_{s'_\ell \in B_\ell^n} \binom{\overline{S}_\ell^n}{s'_\ell} \exp(\beta_\ell^n s'_\ell)},$$

Taking the product of independent binomials,

$$\mathrm{P}\left(\widetilde{S}^n = s | s \in B^n\right) = \prod_\ell \frac{\binom{\overline{S}_\ell^n}{s_\ell} \exp\left(\beta_\ell^n s_\ell\right)}{\sum_{s'_\ell \in B_\ell^n} \binom{\overline{S}_\ell^n}{s'_\ell} \exp\left(\beta_\ell^n s'_\ell\right)}.$$

For a binomial distribution, the probability that $\widetilde{S}_\ell^n \in B_\ell^n \to 1$. Thus, under independent binomial distributions, the probability that $\widetilde{S}^n \in B^n \to 1$, and so the unconditional probability for any $s \in B^n$ converges to

$$\mathrm{P}\left(\widetilde{S}^n = s\right) = \prod_\ell \frac{\binom{\overline{S}_\ell^n}{s_\ell} \exp\left(\beta_\ell^n s_\ell\right)}{\sum_{s'_\ell \in B_\ell^n} \binom{\overline{S}_\ell^n}{s'_\ell} \exp\left(\beta_\ell^n s'_\ell\right)}. \tag{A.3}$$

(A.3) together with (A.2) then implies that the distribution of the SERGM is approximated by the corresponding product of independent binomials.

The remainder of the claimed results then follows easily from standard properties of the binomial distribution. The $V_{\ell\ell}$ terms are computed observing the fact that for $g(x)$ the logistic function,

$$g'(p) = \frac{\exp(p)}{(1+\exp(p))^2}.$$

and one can apply the delta method to obtain the desired result. ■

Proof of Theorem 2.

We begin the proof by showing the following. For any ℓ, the fraction of counts of subnetworks ℓ generated incidentally by some other subnetworks goes to 0. That is, consider some ℓ and $g_\ell \in G_\ell^n$ on m_ℓ nodes. Let

$$\widetilde{p}_\ell^n = \frac{\mathrm{E}(S_\ell^n)}{y_\ell^n \binom{n}{m_\ell}}.$$

This is no more than p_ℓ^n, as the denominator includes all possible subgraphs of size ℓ (where y_ℓ^n is the number of subgraphs of type ℓ that can be formed on any given m_ℓ nodes).[55] Let us consider the probability z_ℓ^n that g_ℓ is incidentally generated by other subnetworks.

We show that $z_\ell^n / \widetilde{p}_\ell^n \to 0$, which implies that $z_\ell^n / p_\ell^n \to 0$.

Consider $g_\ell \in G_\ell^n$ and an incidentally generating subclass $(\ell_j, k_j)_{j \in J}$.

We show that the probability $z_\ell^n(J)$ that it is generated by this subclass goes to zero relative to $\widetilde{p}_\ell^n$, so that $z_\ell^n(J)/\widetilde{p}_\ell^n \to 0$ for each J, and since there are at most $M_\ell \leq k^{m_\ell}$ such generating classes, this implies that $z_\ell^n / \widetilde{p}_\ell^n \to 0$.

[55]Here we provide the proof that applies without subgraphs being characteristic dependent. The extension to characteristic dependent subgraphs is straightforward simply by adjusting all numbers to reflect possible networks with given node characteristics as dependent on n and the sets of nodes that have particular characteristics as a function of n, but it is notationally much more intensive.

For a subnetwork in $G^n_{\ell_j}$, the probability of getting *at least one* such network that has the k_j nodes out of the m_ℓ in g_ℓ is no more than[56]

$$y^n_{\ell_j}\binom{n}{m_{\ell_j}-k_j}\widetilde{p}^n_{\ell_j} \leq y^n_{\ell_j} n^{m_{\ell_j}-k_j}\widetilde{p}^n_{\ell_j}.$$

Thus,

$$\frac{z^n_\ell(J)}{\widetilde{p}^n_\ell} \leq \frac{\Pi_{j\in J} n^{m_{\ell_j}-k_j} y^n_j \widetilde{p}^n_{\ell_j}}{\widetilde{p}^n_\ell}.$$

Therefore

$$\frac{z^n_\ell(J)}{\widetilde{p}^n_\ell} \leq \frac{\Pi_{j\in J} y^n_j n^{m_{\ell_j}} \widetilde{p}^n_{\ell_j}}{n^{\sum_j k_j}\widetilde{p}^n_\ell}.$$

Recall that $M_J = \sum_{j\in J} k_j - m_\ell$ and that $M \geq 1$ (since $|J| \geq 2$ and each set of k_j intersects with at least one other set of $k_{j'}$ for some $j' \neq j$).

Therefore

$$\frac{z^n_\ell(J)}{\widetilde{p}^n_\ell} \leq \frac{\Pi_{j\in J} y^n_{\ell_j} n^{m_{\ell_j}} \widetilde{p}^n_{\ell_j}}{n^{M_J} n^{m_\ell}\widetilde{p}^n_\ell}.$$

The numerator is of the order $\Pi_{j\in J}\mathrm{E}(S^n_j)$ while the denominator is of the order $n^{M_J}\mathrm{E}(S^n_\ell)$.

Under the sparseness condition,

$$\frac{\Pi_{j\in J}\mathrm{E}(S^n_{\ell_j})}{n^{M_J}\mathrm{E}(S^n_\ell)} \to 0,$$

and so we have verified the claim.

To finish the ratio consistency proof, note that the claim then implies that $\frac{\widehat{S}^n_\ell(g)}{S^n_\ell} \to 1$. Thus, dividing top and bottom by $\overline{S}^n_\ell(g)$, it follows that $\frac{\widehat{p}^n_\ell(g)}{S^n_\ell/\overline{S}^n_\ell(g)} \to 1$. Given the growing condition and properties of the binomial distribution, it also follows that $\frac{S^n_\ell/\overline{S}^n_\ell(g)}{p^n_\ell} \to 1$, and so $\frac{\widehat{p}^n_\ell(g)}{p^n_\ell} \to 1$.

Next, note that the above also implies that the distribution $D_n^{1/2}\left(\widehat{p}^n_1(g),...,\widehat{p}^n_k(g)\right)'$ converges to the distribution of $D_n^{1/2}\left(\widetilde{p}^n_1(g),...,\widetilde{p}^n_k(g)\right)'$, where $\widetilde{p}^n_\ell(g) = S^n_\ell/\overline{S}^n_\ell(g)$.

The asymptotic normality of the (joint) distribution then follows from the usual Linderberg-Feller central limit theorem applied to triangular arrays of binomial random variables. This applies under the growing conditions of the theorem.[57] ■

[56]This is a loose upper bound as it simply adds the probability that each possible one forms - but becomes more accurate as the probability of any one occurring vanishes.

[57] Let $X_{1T},...,X_{TT}$ be a triangular array of $\mathrm{Ber}(p_T)$ random variables and $p_T T \to \infty$ (as under the growing condition). To apply the Lindeberg-Feller central limit theorem for triangular arrays, we check the Lindeberg condition: for any $\epsilon > 0$,

$$\frac{1}{Tp_T(1-p_T)}\sum_{t\leq T}\mathrm{E}\left[X_t^2\mathbf{1}\left\{|X_t| \geq \epsilon\sqrt{Tp_T(1-p_T)}\right\}\right] = o(1).$$

The condition is implied by the fact that $Tp_T \to \infty$, and from the sparsity conditions which imply that p_T is bounded away from 1.

Appendix B. Additional Consistency Results

B.1. **Ratio Convergence Results.** A sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$ is *ratio-expectations-identified* with respect to a sequence of diagonal C^n with positive diagonal entries if there exists $\gamma > 0$ such that

$$\left|\frac{C^n_{hh}\mathrm{E}_{\beta}[S^n_h]}{C^n_{hh}\mathrm{E}_{\beta^n}[S^n_h]} - 1\right| > \gamma \left|\frac{\beta_h}{(\beta^n)_h} - 1\right|$$

for all n and h.

A sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$ is *ratio-concentrated* with respect to a sequence of diagonal C^n with positive diagonal entries if

$$\frac{C^n_{hh}S^n_h}{C^n_{hh}\mathrm{E}_{\beta_0}[S^n_h]} - 1 \xrightarrow{\mathrm{P}} 0$$

for each h.

Proposition **B.1.** *If a sequence of SERGMs $(S^n, K^n_{S^n}, A_n, \beta^n)$ is ratio-expectations-identified and ratio-concentrated with respect to a sequence of diagonal C^n with positive diagonal entries, then $\frac{\widehat{\beta}^n(S^n)_h}{(\beta^n)_h} \xrightarrow{\mathrm{P}} 1$ for each h.*

Proof of Proposition B.1. Again, recalling that the MLE $\widehat{\beta}^n(s)$ is the β that solves

$$s = \mathrm{E}_{\beta}[S^n].$$

Thus, since ratio-concentration implies that

$$\frac{C^n_{hh}S^n_h}{\mathrm{E}_{C^n_{hh}\beta^n}[S^n_h]} - 1 \xrightarrow{\mathrm{P}} 0$$

it follows that

$$\frac{C^n_{hh}\mathrm{E}_{\widehat{\beta}^n(S^n)}[S^n_h]}{C^n_{hh}\mathrm{E}_{\beta^n}[S^n_h]} - 1 \xrightarrow{\mathrm{P}} 0.$$

Given that ratio expectations identification implies that

$$\left|\frac{C^n_{hh}\mathrm{E}_{\beta}[S^n_h]}{C^n_{hh}\mathrm{E}_{\beta^n}[S^n_h]} - 1\right| > \gamma\left|\frac{\beta_h}{(\beta^n)_h} - 1\right|$$

for all n, h, it follows that

$$\frac{\widehat{\beta}^n(S^n)_h}{(\beta^n)_h} - 1 \xrightarrow{\mathrm{P}} 0$$

or

$$\frac{\widehat{\beta}^n(S^n)_h}{(\beta^n)_h} \xrightarrow{\mathrm{P}} 1,$$

as claimed. ■

B.2. **Relationship with Binomial Models.**

We now discuss some simple conditions for consistency from the perspective of a binomial distribution. Consider the probability distribution defined over $\left\{0,...,\bar{X}_n\right\}$ by

$$\mathrm{P}_\beta\left(X_n = x\right) = \frac{K_n\left(x\right)\cdot\exp\left(\beta\cdot x\right)}{\sum_{x'} K_n\left(x'\right)\cdot\exp\left(\beta\cdot x'\right)}.$$

We can ask under what assumptions on $\left\{K_n\left(\cdot\right)\right\}_{n\in\mathbb{N}}$ does $\widehat{\beta}\xrightarrow{\mathrm{P}}\beta$? It may seem odd first that a single draw along a sequence of distributions can generate a consistent estimation. However, we make a simple observation that we can transform this distribution into one that resembles a reweighted binomial distribution:

$$\mathrm{P}_\beta\left(X_n = x\right) = \frac{\binom{\bar{X}_n}{x}\cdot\exp\left(\beta\cdot x+\lambda_n\left(x\right)\right)}{\sum_{x'}\binom{\bar{X}_n}{x'}\cdot\exp\left(\beta\cdot x'+\lambda_n\left(x'\right)\right)}, \tag{B.1}$$

where $\lambda_n(x) := \log\left(\omega_n(x)\right) := \log\left(K_n(x)/\binom{\bar{X}}{x}\right)$.

This comes from the fact that $\mathrm{P}_\beta\left(X_n = x\right) = \frac{\binom{\bar{X}_n}{x}\cdot\exp(\beta\cdot x)}{\sum_{x'}\binom{\bar{X}_n}{x'}\cdot\exp(\beta\cdot x')}$ is the distribution corresponding to the probability distribution of a binomial, $\mathrm{Bin}\left(\bar{X}_n; p=\frac{\exp\beta}{1+\exp\beta}\right)$. In turn, assumptions on $\lambda_n\left(x\right)$ will allow for consistency of the MLE of β.

LEMMA **B.1.** *$\widehat{\beta} := \frac{\exp\frac{s_n}{\bar{S}_n}}{1+\exp\frac{s_n}{\bar{S}_n}}$ be the MLE in the above model. Assume $\omega_n\left(\cdot\right)$ is such that*

$$\mathrm{cov}^{\mathrm{Bin}\left(\bar{X}_n;p\right)}\left(x,\omega_n\left(x\right)\right) = o\left(\mathrm{E}^{\mathrm{Bin}\left(\bar{X}_n;p\right)}\left[\omega_n\left(x\right)\right]\right)$$

uniformly in a compact neighborhood of p. Then $\widehat{\beta}\xrightarrow{\mathrm{P}}\beta$.

Proof. First the normalizing constant can be written as

$$\phi\left(\beta\right) = \frac{\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]}{\left(1-p\right)^{\bar{S}_n}}.$$

This follows from

$$\binom{\bar{S}_n}{s_n}\cdot\exp\left(\log\left(\frac{p}{1-p}\right)\cdot s_n+\lambda_n\left(s_n\right)\right) = \binom{\bar{S}_n}{s_n}\cdot\left(\frac{p}{1-p}\right)^{s_n}\times\omega_n\left(s_n\right)$$

and therefore

$$\sum_{s_n'}\binom{\bar{S}_n}{s_n'}\cdot\exp\left(\log\left(\frac{p}{1-p}\right)\cdot s_n'+\lambda_n\left(s_n'\right)\right) = \sum_{s_n'}\binom{\bar{S}_n}{s_n'}\cdot\left(\frac{p}{1-p}\right)^{s_n'}\times\omega_n\left(s_n'\right) = \frac{\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]}{\left(1-p\right)^{\bar{S}_n}}.$$

Second, the maximum likelihood estimator solves the FOC of

$$s_n\log\left(\frac{p}{1-p}\right)+\log\left(1-p\right)^{\bar{S}_n}-\log\left(\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]\right),$$

given by

$$0 = s_n\frac{p}{p\left(1-p\right)}+s_n\frac{1-p}{p\left(1-p\right)}-\frac{\bar{S}_n}{1-p}-\frac{\partial_p\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]}{\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]}.$$

The parameter $\widehat{p}$ is therefore given implicitly by

$$\widehat{p} = \frac{s_n}{\bar{S}_n} - \frac{\partial_p \mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]}{\mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right]} \times \frac{1}{1-\widehat{p}}.$$

Third, observe that

$$\partial_p \mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right] = \frac{1}{p\left(1-p\right)} \mathrm{E}_p^{Bin}\left[\left\{s_n' - p\bar{S}_n\right\}\omega_n\left(s_n'\right)\right],$$

which follows from

$$\begin{aligned}
\partial_p \mathrm{E}_p^{Bin}\left[\omega_n\left(s_n'\right)\right] &= \sum_{s_n'} \binom{\bar{S}_n}{s_n'} \cdot \partial_p \left\{p^{s_n'}\left(1-p\right)^{\bar{S}_n - s_n'}\right\} \times \omega_n\left(s_n'\right) \\
&= \sum_{s_n'} \binom{\bar{S}_n}{s_n'} \cdot p^{s_n'}\left(1-p\right)^{\bar{S}_n - s_n'} \times \omega_n\left(s_n'\right) \times \left\{\frac{s_n'}{p} - \frac{\left(\bar{S}_n - s_n'\right)}{1-p}\right\} \\
&= \frac{1}{p\left(1-p\right)} \mathrm{E}_p^{Bin}\left[\left\{s_n' - p\bar{S}_n\right\}\omega_n\left(s_n'\right)\right].
\end{aligned}$$

Thus

$$\widehat{p} = \frac{s_n}{\bar{S}_n} - \frac{\mathrm{E}_{\widehat{p}}^{Bin}\left[\left\{s_n' - p\bar{S}_n\right\}\omega_n\left(s_n'\right)\right]}{\mathrm{E}_{\widehat{p}}^{Bin}\left[\omega_n\left(s_n'\right)\right]} \times \frac{1}{1-\widehat{p}}.$$

Under the assumed condition that $\sup_{p' \in B_\delta(p)} \frac{\mathrm{cov}^{Bin(\bar{S}_n;p')}(s,\omega_n(s))}{\mathrm{E}^{Bin(\bar{S}_n;p')}[\omega_n(s)]} = o\left(1\right)$, the result follows. ■

B.3. **A Characterization of Consistency in terms of Variance.** We say that a SERGM is *expectations-identified* if $\beta' \neq \beta$ implies that $\mathrm{E}_\beta[S] \neq \mathrm{E}_{\beta'}[S]$.

For simplicity, we present the univariate case, though the multivariate case follows simply by controlling the norm of the covariance matrix instead.

PROPOSITION **B.2.** *Consider an expectations-identified SERGM.*

(1) *For any* $\varepsilon > 0$ *such that* $\varepsilon \leq \frac{\mathrm{var}_\beta[S]}{\max_{\beta' \in [\beta-\varepsilon,\beta+\varepsilon]} |k^3_{\beta'}(S)|}$,

$$\mathrm{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \leq \frac{4}{\varepsilon^2 \mathrm{var}_\beta(S)}.$$

(2) *Conversely, let* $B > 0$ *be such that* $\mathrm{var}_\beta[S] \leq B$ *and* $|S - \mathrm{E}_{\beta'}[S]| \leq B\mathrm{E}_{\beta'}[|S - \mathrm{E}_{\beta'}[S]|]$ *for all* β' *within* $\frac{1}{(B+1)B}$ *of* β. *Then*

$$\mathrm{P}_\beta^n\left[|\widehat{\beta}^n(S) - \beta| > \frac{1}{3(B+1)B}\right] > \frac{1}{2B-1}.$$

Consider a sequence of expectations-identified SERGMs indexed by the number of nodes n. We say that the sequence is consistent at β if there are sequences $\varepsilon^n \to 0$ and $r^n \to \infty$ such that

$$\mathrm{P}_\beta^n\left[|\widehat{\beta}^n(S) - \beta| > \varepsilon^n\right] < \frac{1}{r^n}.$$

COROLLARY **B.1.** *Consider a sequence of expectations-identified SERGMs indexed by the number of nodes* n.

(1) *If* $\mathrm{var}^n_\beta[S] \to \infty$ *and* $\mathrm{Skew}^n_\beta[S] \to 0$*, then the sequence is consistent at* β*.*
(2) *Conversely, if there is some* $B > 0$ *such that* $\mathrm{var}^n_\beta[S] \leq B$ *and* $|S - \mathrm{E}^n_{\beta'}[S]| \leq B\mathrm{E}^n_{\beta'}\left[\left|S - \mathrm{E}^n_{\beta'}[S]\right|\right]$ *for all* β' *within* $\frac{1}{(B+1)B}$ *of* β *for large enough* n*, then the sequence is not consistent.*

To see an implication of this result, consider a SERGM defined on a count of "large" subgraphs, for instance the number of components that contain some given percentage of the nodes. The variance of such a statistic is necessarily bounded, and so such a SERGM cannot be consistent.

Proof of Proposition B.2.

First, note that since $\beta' \neq \beta$ implies that $\mathrm{E}_\beta[S] \neq \mathrm{E}_{\beta'}[S]$, and also since $\mathrm{E}_\beta[S]$ is continuous in β, it must be that $\mathrm{E}_\beta[S]$ is monotone in β.

Let $\widehat{\beta}(S)$ be the MLE of β as a function of the statistic S.

Note that from Section 3.4 we know that $\widehat{\beta}(s)$ is the β such that $E_\beta[S] = s$, which varies continuously in β.

Let us consider

$$\mathrm{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right]$$

Note that by Lemma B.2, it follows that $s'(\beta) = \mathrm{var}_\beta(S)$, and also that $s''(\beta) = \mathrm{E}_\beta\left[(S - \mathrm{E}_\beta[S])^3\right]$.

Therefore,

$$s(\beta + \varepsilon) = s(\beta) + \int_0^\varepsilon \left[\mathrm{var}_\beta(S) + \int_0^x k^3_{\beta+y}(S)dy\right] dx$$

and

$$s(\beta - \varepsilon) = s(\beta) - \int_0^\varepsilon \left[\mathrm{var}_\beta(S) + \int_0^x k^3_{\beta-y}(S)dy\right] dx.$$

Thus,

$$\min\left[|s(\beta+\varepsilon) - s(\beta)|, |s(\beta-\varepsilon) - s(\beta)|\right] \geq \varepsilon\mathrm{var}_\beta(S) - \frac{\varepsilon^2}{2}\left(\max_{\beta' \in [\beta-\varepsilon, \beta+\varepsilon]} |k^3_{\beta'}(S)|\right),$$

and by the choice of $\varepsilon \leq \frac{\mathrm{var}_\beta[S]}{\max_{\beta' \in [\beta-\varepsilon,\beta+\varepsilon]} |k^3_{\beta'}(S)|}$,

$$\min\left[|s(\beta+\varepsilon) - s(\beta)|, |s(\beta-\varepsilon) - s(\beta)|\right] \geq \varepsilon\mathrm{var}_\beta(S)/2.$$

Given the monotonicity of $E_\beta(S)$ in β, it then follows that

$$|\widehat{\beta}(S) - \beta| > \varepsilon \text{ implies } |S - \mathrm{E}_\beta(S)| > \varepsilon\mathrm{var}_\beta(S)/2.$$

Thus,

$$\mathrm{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \leq \mathrm{P}_\beta\left[|S - \mathrm{E}^n_\beta(S)| > \varepsilon\mathrm{var}_\beta(S)/2\right].$$

By Chebychev's inequality it then follows that

$$\mathrm{P}_\beta\left[|S - \mathrm{E}^n_\beta(S)| > \varepsilon\mathrm{var}_\beta(S)/2\right] \leq \frac{\mathrm{var}_\beta(S)}{(\varepsilon\mathrm{Var}_\beta(\mathrm{S})/2)^2} = \frac{4}{\varepsilon^2\mathrm{Var}_\beta(\mathrm{S})}.$$

Therefore,

$$\mathrm{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \leq \frac{4}{\varepsilon^2\mathrm{Var}_\beta(\mathrm{S})},$$

as claimed.

Now let us examine the converse statement in the proposition.

Note that $\text{var}_\beta[S] \leq B$, it follows that $\text{E}[|S - \text{E}^n_\beta(S)|] \leq B + 1$, and therefore by assumption it also follows that $|S - \text{E}^n_\beta(S)| \leq (B + 1)B$.[58] It then follows that

$$\text{var}^n_\beta(S) \leq |S - \text{E}^n_\beta(S)|\text{E}[|S - \text{E}^n_\beta(S)|] \leq (B + 1)B\text{E}[|S - \text{E}^n_\beta(S)|],$$

and by similar reasoning

$$k^3_\beta[S] \leq |S - \text{E}^n_\beta(S)|\text{var}^n_\beta(S) \leq (B + 1)^2B^2\text{E}[|S - E^n_\beta(S)|].$$

By similar reasoning to the proof of the first part of the proposition, we can deduce that

$$\text{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \geq \text{P}_\beta\left[|S - E_\beta(S)| > \varepsilon\text{Var}_\beta(\text{S}) + \varepsilon^2\left(\max_{\beta' \in [\beta-\varepsilon,\beta+\varepsilon]} |\text{k}^3_{\beta'}(\text{S})|\right)\right].$$

Therefore,

$$\text{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \geq \text{P}_\beta\left[|S - E^n_\beta(S)| > \left(\varepsilon(B + 1)B + \varepsilon^2(B + 1)^2B^2\right) E[|S - \text{E}^n_\beta(S)|]\right].$$

For $\varepsilon = \frac{1}{3(B+1)B}$, it follows that

$$\text{P}_\beta\left[|\widehat{\beta}(S) - \beta| > \varepsilon\right] \geq \text{P}_\beta\left[|S - \text{E}^n_\beta(S)| > \text{E}[|S - \text{E}^n_\beta(S)|]/2\right].$$

Given that $|S - \text{E}^n_\beta(S)| < B\text{E}[|S - \text{E}^n_\beta(S)|]$, it follows that there exists $\delta = \frac{1}{2B-1}$ such that

$$\text{P}_\beta\left[|S - \text{E}^n_\beta(S)| > \text{E}[|S - \text{E}^n_\beta(S)|]/2\right] \geq \delta.$$

establishing the second part of the proposition.

To see this last claim, let $p = \text{P}_\beta\left[|S - \text{E}^n_\beta(S)| > \text{E}[|S - \text{E}^n_\beta(S)|]/2\right]$ and note that

$$(1-p)\text{E}\left[|S - \text{E}^n_\beta(S)| \;\middle|\; |S - E^n_\beta(S)| \leq \text{E}[|S - \text{E}^n_\beta(S)|]/2\right] + p\text{E}\left[|S - \text{E}^n_\beta(S)| \;\middle|\; |S - \text{E}^n_\beta(S)| > \text{E}[|S - \text{E}^n_\beta(\right.$$

$$= \text{E}[|S - \text{E}^n_\beta(S)|]$$

and so

$$(1 - p)\text{E}[|S - \text{E}^n_\beta(S)|]/2 + pB\text{E}[|S - \text{E}^n_\beta(S)|] \geq \text{E}[|S - \text{E}^n_\beta(S)|].$$

Therefore, $p \geq \delta = \frac{1}{2B-1}$, as claimed. ■

Proof of Corollary B.1.

The second claim follows directly from Proposition B.2.

To see the first claim, let us consider two cases. First consider a case such that $Var^n_\beta(S)/|k^3_\beta(S)^n|$ is bounded away from 0. In that case apply the first part of Proposition B.2, with $\varepsilon^n = (Var^n_\beta(S))^{-a}$ for any $a < 1/2$. Next consider a case such that $\text{var}^n_\beta(S)/|k^3_\beta(S)^n| \to 0$. In that case, set $\varepsilon^n = \frac{\text{var}^n_\beta(S)}{|k^3_\beta(S)^n|}$. In that case, we need to check that $\frac{(\text{var}^n_\beta(S))^3}{|k^3_\beta(S)^n|^2} \to \infty$, which is equivalent to $\text{Skew}^\text{n}_\beta(\text{S}) \to 0$. ■

[58] To see that $\text{E}[|S - \text{E}^n_\beta(S)|] \leq B + 1$, note that for a nonnegative random variable X,

$$\text{E}[X] \leq \text{P}[X \leq 1]1 + \text{P}[X > 1]\text{E}[X|X > 1]$$

and

$$\text{P}[X > 1]E[X|X > 1] < \text{P}[X > 1]\text{E}[X^2|X > 1] \leq \text{E}[X^2],$$

and the claim follows.

The exponential family gives us a useful relationship between s and its cumulants that is useful in what follows.

LEMMA **B.2** (Cumulant expansion)**.** *$s^{(q)}(\beta)$ is the $(q-1)$th cumulant of S.*

Proof of Lemma B.2. This follows by inspection. For exponential families $\kappa_S(t) = \log\left(\frac{\phi(\beta+t)}{\phi(\beta)}\right)$. Then the fact that the moment generating function can be written as $M_S(t) = \sum_s \frac{K_S(s)}{\phi(\beta)} e^{\beta \cdot s} e^{t \cdot s} = \frac{\phi(\beta+t)}{\phi(\beta)}$ and the cumulant generating function is the log of the moment generating function implies the result. ■